The Shnoudi Descent: From the Golay Code and Leech Lattice to Four-Dimensional Triality — 24 to 4

Version Two — A Speculative Mathematical Framework for the Emergence of Four-Dimensional Geometry and Particle Symmetry

Ibrahim Amin Nicola Shnoudi — Independent Researcher —
Ibrahim.shnoudiest@gmail.com

Citation and use: quotation is permitted provided both the document's name and the author's name are cited.

Parent theory: The Shnoudi Admissibility Theory (SAT), Edition 2025–2026 — DOI: 10.5281/zenodo.18153991. The complete code appendix (55 scripts) is provided as a separate companion file, archived alongside this document on Zenodo under the same DOI.

Abstract

This document develops a mathematical construction connecting the Golay code and Leech lattice, in twenty-four dimensions, to a four-dimensional, triality-based particle symmetry structure. Every computational claim is verified by 55 independently runnable Python scripts, provided as a separate companion code file archived alongside this document. This version's central finding: the triality automorphism used throughout Parts VII to IX is shown, by explicit construction rather than general argument, to be inherited from the Mathieu group M_{24} — the symmetry group of the Golay code this document starts from — rather than imposed on D_4 independently. A genuine, verified order-3 element of M_{24} is shown, via a fully explicit chain of orthogonal matrices, to induce this document's own triality generator exactly, confirmed with zero exceptions across every independent test performed (Section 24b). Two companion clarifications close the document's two central selection questions as completely as their respective hypotheses allow: Section 23.4 states the single physical hypothesis — an order-3 outer automorphism, motivated by three fermion generations — under which D_4's selection among all simple Lie algebras is a complete logical consequence, and Section 3.2 states the analogous hypothesis — lattice minimality combined with well-definedness — under which dimension 24 and the Leech lattice are shown to be the unique admissible starting point among every classified dimension. Both proofs are complete conditional on their stated hypothesis; neither hypothesis is itself derived, and their differing epistemic character — one anchored to a direct experimental count, the other to a methodological principle — is stated explicitly rather than elided. Whether the M_{24} connection has been noted elsewhere — it follows fairly directly from classical facts (E_8's uniqueness, Borel–de Siebenthal conjugacy) well known to specialists in this area, and may be folklore even where not written up — was not established by a systematic literature search and is not claimed here; only that this document did not previously contain it, and now does, with full derivation and verification. This document distinguishes throughout, and states plainly, which results are exact mathematical or computational facts and which remain open physical questions; as of this version, it offers no experimentally testable prediction.

Table of Contents

Introduction 4

Notation and Conventions 4

Part I — Foundations: The Golay Code and the Leech Lattice 5

1. The Classical Objects 5
2. A Principled Selection 5
3. The Physical Reading of Dimension 24 5
Part II — The Descent to Four Dimensions 8
4. The Laminated Sequence 8
5. Structural Landmarks 8
Part III — From Discrete Lattice to Continuum Geometry 10
6. Convergence to Euclidean Space 10
7. Intrinsic Geometry 10
Part IV — $E_8 \supset D_4 \oplus D_4$: Explicit Verification 11
8. The Classical Embedding 11
9. Direct Computational Verification 11
Part V — From E_8 to Standard Model Quantum Numbers 12
10. The $E_6 \times SU(3)$ Breaking 12
11. The Extended Chain to One Generation 12
Part VI — Independent Physical Corroboration 14
12. The Conway Group in Heterotic String Compactifications 14
13. An Explicit Verified Isometry 14
Part VII — Toward Three Triality Sectors: A New Cyclic Structure 15
14. A Triality-Invariant SU(3), Found and Verified 15
15. The Question 17
16. Three Cyclically Related Copies, Verified 17
17. The Combinatorial Structure of Overlap, Fully Resolved 18
18. Scope and Limitations of This Result 21
Part VIII — Toward Dynamics and Chirality 22
19. The Chirality Obstruction 22
20. A Framework for Dynamics 34

Part IX — The Search for a Non-Circular Selection Principle 36
21. Criteria Tested and Their Outcomes 36
22. A Clean, Non-Circular Result 36
23. Necessary and Sufficient Conditions Within the Classification of Simple Lie Algebras 37
24. Relation to Existing Approaches 40
24b. Triality Is Inherited from M_{24}, Not Imposed on D_4 Independently 42
24c. A Genuine S_3, Not Merely Its Order-3 Subgroup 44
24d. The Full Trio: Six Copies of T, Not One 45
24e. A Second, Structurally Distinct Order-3 Symmetry in M_{24} 46
Original Contributions of This Work 48
Part X — Status 52
25. What Is Established, What Is Postulated, What Remains Open 52
Appendix — Computational Reproducibility 57
References 60

A Discrete Descent from the Golay Code and Leech Lattice to D_4, E_6, and Three Cyclic Triality Sectors

Introduction

Between a single point and the continuum lies every structure mathematics has built to describe space. This document follows a deliberately chosen thread through that territory: a path descending from a twenty-four-dimensional lattice — the Leech lattice — through a sequence of smaller, related structures, arriving at dimension four with an explicit account of how continuous four-dimensional geometry, its internal symmetry structure, and a viable route toward Standard Model quantum numbers and generation structure can emerge from discrete, combinatorial beginnings.

This is a speculative theoretical-physics construction, not a claim about the confirmed structure of the universe. Every mathematical and computational step below is independently checkable and is presented as such; every point at which the argument depends on a chosen hypothesis rather than a proof is identified plainly. None of the individual mathematical ingredients assembled here — the Leech lattice, E_8's root system, D_4 triality, the E_6/SO(10) embedding, the Candelas–Horowitz–Strominger–Witten (CHSW) breaking — is new; what is offered as this document's contribution is their synthesis into one connected, computationally verified chain, not a claim of newly discovered mathematics.

Golay Code → Leech Lattice (Λ_{24}) → Λ_{16} → E_8 ⊃ $D_4 \oplus D_4$ → $E_6 \times SU(3)$ → [E_7 → 27 → SO(10), one generation] → Continuum Limit (D_{4a}) → Spin(8) Triality → Three Cyclic Triality Sectors

The pathway this document develops, in overview.

Notation and Conventions

$\Phi(G)$ denotes the root system of a Lie algebra or lattice G; W(G) its Weyl group; $\Delta(G)$ its Dynkin diagram. A_n, D_n, E_n refer to the standard Cartan–Killing classification of simple Lie types. Λ_m denotes the laminated lattice of rank m (Part II); D_4 without further qualifier denotes the rank-4 checkerboard lattice of Part III. D_{4a} and D_4_B denote the two orthogonal rank-4 factors of the $E_8 \supset D_4 \oplus D_4$ decomposition of Part IV, labeled "spacetime" and "internal" respectively as an external convention, not a distinction forced by the mathematics (Part V, Section 10). T denotes the order-three triality automorphism constructed and verified in Part VII; 8_v, 8_s, 8_c denote the vector and two half-spin representations of D_4 permuted by T. All roots are normalized to norm^2 = 2 throughout, following the standard simply-laced convention; inner products are the standard Euclidean pairing on the ambient coordinate space stated at each construction's introduction ($\mathbb{R}^8$ for E_8, $\mathbb{R}^4$ for D_4). Coordinate conventions for E_8's 240 roots and D_4's 24 roots are stated explicitly where first used (Part IV, Section 9; Part III, Section 6) and used consistently thereafter.

Part I — Foundations: The Golay Code and the Leech Lattice

1. The Classical Objects

The binary Golay code $\mathcal{G}_{24}$ is the [24,12,8] linear code over $\mathbb{F}_2$: a 12-dimensional subspace of $\mathbb{F}_2{}^{24}$ with minimum Hamming distance 8, self-dual ($\mathcal{G}_{24} = \mathcal{G}_{24}\perp$), automorphism group the sporadic simple Mathieu group M_{24} (order 244,823,040). An explicit realization via the bordered quadratic-residue construction on eleven points, with generator matrix $G = [I_{12} \mid B]$ built from quadratic residues modulo 11, was verified directly: all 4,096 codewords enumerate to the exact weight distribution {0:1, 8:759, 12:2576, 16:759, 24:1}, and self-orthogonality $G \cdot G^T = 0 \pmod 2$ holds exactly.

The Leech lattice Λ_{24} is the unique 24-dimensional even, unimodular, rootless lattice, constructed explicitly from the Golay code via Conway's Miracle Octad Generator (see also Borcherds, 1984/1987, for the classical thesis-length treatment of this and related lattice constructions). Its automorphism group Co_0 has order 8,315,553,613,086,720,000; $Co_0/\{\pm 1\}$ is the sporadic simple group Co_1.

2. A Principled Selection

Among the 24 even unimodular lattices in dimension 24 (the Niemeier lattices), Leech is selected by three explicit criteria — even, unimodular, rootless — stated before checking what they select. Leech is the unique Niemeier lattice satisfying rootlessness; every alternative — including $E_8 \oplus E_8 \oplus E_8$, the Niemeier lattice with root system A24 (Niemeier(A24)), and the Niemeier lattice with root system D24 (Niemeier(D24)) — contains a nonempty root system and is excluded directly. By the Frenkel–Kac–Segal mechanism (Part II, Section 5), a nonempty root system generates non-abelian gauge symmetry immediately at the starting dimension; only Leech, among all 24 Niemeier lattices, defers the appearance of such symmetry to a later stage of the descent, giving the rootlessness criterion a structurally distinguished reading beyond its role as a uniqueness device.

3. The Physical Reading of Dimension 24

Two independent characterizations converge on this dimension. First, a direct classification fact: even unimodular lattices exist only in dimensions divisible by eight, and at each of the two smaller candidates rootlessness is impossible — dimension 8 has a unique even unimodular lattice (E_8 itself, 240 roots), and dimension 16 has exactly two ($E_8 \oplus E_8$ and $D_{16}{}^+$), both of which contain roots inherited from their D_{16} sublattice structure. Twenty-four is therefore the first dimension where rootlessness is even possible, independent of any further argument. A classical modular-form argument sharpens this further: 24 is the smallest rank at which an even unimodular lattice can be rootless as a consequence of weight-12 modular forms admitting a cusp form, and the resulting rootless lattice is canonical because the relevant space of cusp forms is one-dimensional. Second, an independent vertex-operator-algebra characterization: Schellekens (1993) classified all 71 holomorphic conformal field theories of central charge 24, and the Leech (Monster) module is the unique case with vanishing weight-one space — the vertex-operator-algebra mirror

of rootlessness, read as the unique such theory carrying no continuous internal gauge symmetry. The uniqueness of this module among all c = 24 holomorphic vertex operator algebras with vanishing weight-one space remains a conjecture in the literature (Frenkel–Lepowsky–Meurman, 1988), named here as the precise residual gap in this second characterization, not concealed.

3.1. Why Not a Higher Dimension: The Loss of Uniqueness Beyond 24

The preceding section establishes that no dimension below 24 admits a rootless even unimodular lattice at all. The complementary question — why not a dimension above 24, where rootless examples do exist — requires a separate argument, since rootlessness alone does not disappear at higher dimensions.

It is answered by the same uniqueness property already used to select Leech among the 24 Niemeier lattices, examined at higher dimensions rather than assumed to hold there. At dimension 32, rootless even unimodular lattices exist, but King's classification shows there are more than ten million of them, with no single lattice distinguished as uniquely extremal in the way the Leech lattice is; a complete classification at dimension 48 is considered infeasible outright. The uniqueness that makes Leech the single, unambiguous outcome of the rootlessness criterion is therefore not a generic feature of large dimensions — it holds at dimension 24 specifically and fails immediately beyond it, with the number of rootless competitors growing from one, to millions, within a single further step of eight. A joint criterion combining rootlessness with uniqueness among rootless lattices — expressed independently of the packing-density considerations of Part III and checkable directly from lattice classification data — is accordingly satisfied, among all dimensions classified to date (up to 48), only at dimension 24 (dimension 8 satisfies uniqueness alone, E_8 being the unique even unimodular lattice there, but fails rootlessness, since E_8 itself has 240 roots).

This last point closes a related question directly: E_8 is excluded as a starting dimension for exactly the same reason dimension 16 is excluded, not a different one — its unique even unimodular lattice is maximally rootful, not rootless, so the entire physical motivation for rootlessness (deferring the appearance of gauge symmetry rather than assuming it from the outset) fails at the first possible step. Within this document's own construction, E_8 is not an alternative starting assumption but a derived waypoint of the ongoing descent from Leech, reached at rank 8 and cross-verified by two independently motivated mechanisms (Part II, Section 4) — a status unavailable to it as a starting point in its own right.

3.2. A Complete Conditional Proof: Why 24, Given a Minimality-and-Uniqueness Hypothesis

Sections 3–3.1 established, as a direct classification fact independent of any preference, that dimension 24 is simultaneously the *first* dimension admitting a rootless even unimodular lattice and the *last* dimension (among those classified, to 48) where such a lattice is *unique*. This subsection states the one explicit hypothesis under which that fact closes completely into a uniqueness proof for the starting dimension — in the same conditional form, and with the same care about its limits, as Section 23.4 provides for D_4.

The hypothesis, stated as a hypothesis, not derived here. Suppose nature's fundamental structure is required to be an even unimodular lattice that (a) defers the appearance of non-abelian gauge symmetry beyond the starting point — the physical reading of rootlessness already given via the Frenkel–Kac–Segal mechanism in Section 2 — and (b) is the *lowest-dimensional* such lattice for which this requirement picks out a single, unambiguous candidate, rather than one arbitrary choice among many equally valid ones. Call this Hypothesis L (lattice minimality-and-uniqueness). Condition (b) reflects a standalone methodological principle used elsewhere in physics — that a selection requirement failing to determine a unique outcome cannot, by itself, be the complete criterion — rather than a direct experimental headcount of the kind Hypothesis G invokes for three generations; this difference in character is stated plainly and returned to below.

The theorem, already established. Given Hypothesis L, Sections 3–3.1 apply directly: even unimodular lattices exist only in dimensions divisible by eight; dimension 8 has a unique even unimodular lattice (E_8) but it has roots, failing (a); dimension 16 has exactly two, both with roots, again failing (a); dimension 24 has 24 such lattices, of which exactly one (Leech) is rootless, satisfying both (a) and (b) simultaneously; dimension 32 admits rootless examples but King's classification shows more than ten million of them, failing (b). Among all dimensions classified to date, 24 is the unique dimension satisfying both conditions at once — and it is moreover the *first* dimension satisfying (a) at all, so no smaller candidate is available even before condition (b) is checked.

The conclusion, and its exact logical status. Given Hypothesis L, dimension 24 and the Leech lattice specifically are the unique admissible starting point — not the best of several options, but the only one satisfying the stated requirement among every dimension classified. This is the identical logical form as Section 23.4: a conditional theorem, complete given its antecedent, silent on the antecedent itself.

The honest difference from Hypothesis G, stated directly. Hypothesis G rests on a single, hard, independently measured number: three generations, observed. Hypothesis L rests on a methodological principle — minimality combined with well-definedness — that is common in physics (a theory whose selection criterion admits millions of equally valid outcomes is not usually considered to have identified a mechanism at all) but is not tied to any single independent experimental count in the same direct way. This document does not claim the two hypotheses carry equal evidential weight; it claims only that each, granted on its own terms, closes its respective selection problem completely and non-circularly, and that Hypothesis L is the most precise form in which "why 24" can currently be closed. Whether a hypothesis of Hypothesis G's more direct, empirically-anchored character exists for dimension 24 — waiting to be identified — remains open, exactly as Section 23.4 already states.

Part II — The Descent to Four Dimensions

4. The Laminated Sequence

Conway and Sloane's laminated lattice construction defines, for each dimension m, the densest lattice Λ_m obtainable by stacking copies of Λ_{m-1}, giving a determined descending sequence: Λ_{24} = Leech, Λ_{16} = Barnes–Wall, $\Lambda_8 = E_8$, $\Lambda_6 = E_6$, $\Lambda_5 = D_5$, $\Lambda_4 = D_4$, $\Lambda_1 = \mathbb{Z}$. Over the range where this sequence coincides with the classical root lattices, it is independently reproduced by extending nested root subsystems ($A_1 \subset A_2 \subset \ldots \subset E_8$) — a construction motivated on entirely different grounds. The agreement of two independently motivated mechanisms is evidence the sequence is not arbitrary. Where the laminated construction is non-unique at a given dimension, the branch with the largest automorphism group is selected — a Maximum Symmetry Principle. Symmetry-based selection has substantial precedent in physics without thereby becoming a proof: Noether's theorem ties conservation laws directly to continuous symmetries, gauge theory constructs the Standard Model's dynamics by demanding invariance under prescribed groups, and Klein's Erlangen program organizes geometry itself by which transformation group is preserved. This motivates treating symmetry as a reasonable selection criterion; it does not establish that nature selects the most symmetric option at each branch point, and no such claim is made. This is better named a Maximum Symmetry Selection Hypothesis than a principle: an alternative variational principle, not currently identified, may replace it, and nothing here excludes that nature instead prefers a less symmetric configuration at some branch points, as spontaneous symmetry breaking does in the Standard Model itself. To be explicit: this document does not claim the hypothesis is physically fundamental, only that it is the simplest rule reproducing the stated descent, pending the variational principle named in Part IX. A different selection hypothesis, were one identified, could in principle select a different branch at one or more of these points, leading to a different descent path than the one developed here. Among canonical dimension-reduction procedures on lattices, laminated descent is distinguished by preserving three properties simultaneously — unimodularity, maximal packing density at each step, and compatibility with the Leech lattice as the dimension-24 endpoint. Voronoi cell decomposition and Delone triangulation preserve local geometric data but neither guarantees unimodularity at each intermediate step nor is canonically compatible with Leech's own construction; projection onto a fixed coordinate subspace preserves neither density nor, in general, unimodularity; fixed-sublattice restriction preserves unimodularity only for a sparse, non-canonical set of embedding choices. Laminated descent is not claimed to be the unique procedure satisfying any single one of these properties, but is, among the procedures compared, the only one satisfying all three simultaneously.

5. Structural Landmarks

Two further criteria isolate specific points. Geometric criterion: minimal vectors first reach norm 2 — generating non-abelian symmetry by the Frenkel–Kac–Segal mechanism of lattice vertex operator algebras — at m = 8; Barnes–Wall (16) and Coxeter–Todd (12) are confirmed rootless, E_8 is a root lattice. Algebraic criterion: each Λ_m's Dynkin diagram automorphism group, within the root-lattice range, falls into one of three tiers — trivial

(E_8, E_7), nontrivial abelian (E_6, D_5), or non-abelian (D_4 alone, order six) — the unique member of this family with S_3, not merely $\mathbb{Z}/2$, diagram symmetry. This three-tier classification is the coarsest nontrivial invariant available: any classification of finite groups coarser than trivial/nontrivial collapses to a single class and distinguishes nothing; the next-coarsest refinement beyond trivial/nontrivial is exactly the abelian/non-abelian distinction, since it is the first group-theoretic property not already implied by cardinality alone. Any classification finer than this three-tier scheme requires an additional chosen parameter, which this scheme avoids entirely. These identify three landmarks — 8, 6, 4 — without reference to which was being sought.

Part III — From Discrete Lattice to Continuum Geometry

6. Convergence to Euclidean Space

Let $D_4 := \{(x_1,x_2,x_3,x_4) \in \mathbb{Z}^4 : x_1+x_2+x_3+x_4 \text{ even}\}$. As $N \to \infty$, $((1/N)D_4, d_Euclidean, 0)$ converges to $(\mathbb{R}^4, d_Euclidean, 0)$ in the pointed Gromov–Hausdorff sense: the inclusion map preserves distances exactly, and every full-rank lattice has finite covering radius ρ, giving covering radius $\rho/N \to 0$ for the rescaled family — the density condition required by the ε-approximation formulation, verified including the density-within-bounded-balls technical step.

7. Intrinsic Geometry

A stronger claim — that Euclidean geometry is intrinsic to D_4's own combinatorial structure, not merely inherited from an ambient embedding — was tested against two candidates. Graph (shortest-path) distance fails: on the simpler lattice $\mathbb{Z}^n$ it reduces to the L^1, not L^2, norm, and direct computation on D_4's own 24 minimal-vector generators confirms non-constant distance ratios (0.707 to 1.000) across sample points, ruling out any rescaling that would recover a Euclidean metric.

The diffusive candidate succeeds. For the symmetric random walk stepping to a uniformly random minimal vector, the covariance matrix Σ is invariant under D_4's automorphism group — the Weyl group of an irreducible root system, hence acting irreducibly on the reflection representation — so by the real version of Schur's lemma — the commutant of an irreducible real representation is one-dimensional — Σ is a scalar multiple of the identity; direct computation gives $\Sigma = 0.5 \cdot I$ exactly. The Donsker invariance principle (whose hypotheses — independent, identically distributed, centered, finite-variance increments — hold directly here) then gives convergence to isotropic Brownian motion. The passage from this limiting diffusion to a genuine metric structure proceeds through the Dirichlet form associated with the limiting process: the theory of Dirichlet forms (Fukushima's foundational framework, Cheeger's energy functional on metric measure spaces, and Ambrosio–Gigli–Savaré's analysis of gradient flows in this setting, together with Sturm's and Kigami's treatments of self-similar and limit spaces) shows that, under standard regularity assumptions, an isotropic diffusion determines its generator's Dirichlet form uniquely, the Dirichlet form determines a Laplacian, and an isotropic Laplacian on $\mathbb{R}^4$ determines the flat Riemannian metric tensor up to the single overall scale already left undetermined by the Gromov–Hausdorff limit of Section 6 — completing the chain diffusion to Dirichlet form to Laplacian to metric tensor, with each arrow a standard, cited construction. In the resulting limit, geodesics are straight lines, consistent with Schilder's theorem: for a diffusively rescaled process with isotropic covariance, the large-deviation rate function is the Dirichlet energy $I(\gamma) = \frac{1}{2}\int_0^1 \|\gamma'(t)\|^2 dt$, and the Euler-Lagrange equation for this functional, $\gamma''(t) = 0$, is satisfied exactly by straight-line paths — identifying them as the minimizers, not merely as a plausible reading of the isotropic limit.

Part IV — $E_8 \supset D_4 \oplus D_4$: Explicit Verification

8. The Classical Embedding

Borel–de Siebenthal theory classifies the maximal-rank semisimple subalgebras of E_8: D_8, $A_1 \times E_7$, A_8, $A_2 \times E_6$, $A_4 \times A_4$. $D_4 \oplus D_4$ arises one step further, as the standard block-diagonal decomposition $D_8 = so(16) \supset so(8) \oplus so(8) \cong D_4 \oplus D_4$. Borel–de Siebenthal theory additionally guarantees that any two maximal-rank subalgebras of the same isomorphism type are conjugate under the Weyl group $W(E_8)$, so every embedding of $D_4 \oplus D_4$ in E_8 is $W(E_8)$-conjugate to the one used below; Part VI, Section 13 exhibits this conjugacy explicitly as a verified orthogonal matrix.

9. Direct Computational Verification

The 240 roots of E_8 were generated explicitly — 112 integer-type roots $\pm e_i \pm e_j$ and 128 half-integer-type roots $(\pm\tfrac{1}{2})^8$ of even sign parity — with norm^2 = 2 confirmed throughout. Splitting the eight coordinates into two blocks of four gives two disjoint 24-root sets, each independently confirmed to carry the exact D_4 inner-product distribution $\{\mp 2{:}1, \mp 1{:}8, 0{:}6\}$, and confirmed mutually orthogonal in every pairwise product: an explicit, machine-verified $E_8 \supset D_{4a} \oplus D_{4_}B$. The remaining 192 roots split into three disjoint sets of exactly 64, each confirmed — by checking that projections onto each D_4 factor take 8 distinct values with all 64 combined pairs pairwise distinct — to carry genuine, collision-free (8,8)-bimodule structure, identifying them as the three inequivalent 8-dimensional irreducible representations of D_4: the vector representation 8_v and the two half-spin representations 8_s, 8_c. This confirms the classical Kostant decomposition 248 = 28 + 28 + 64 + 64 + 64 at the level of explicit root vectors. This computation, and the similar finite checks throughout Parts IV through VII (root counts, inner-product distributions, orbit sizes), constitute exhaustive finite verification on explicit, fully enumerated root systems, not symbolic proof in the sense of Parts III and IX's Schur's-lemma, Gromov–Hausdorff, and Necessary-and-Sufficient-Conditions arguments; the two categories are kept distinct throughout, and Part X labels each result accordingly.

Part V — From E_8 to Standard Model Quantum Numbers

10. The $E_6 \times SU(3)$ Breaking

Labeling one D_4 block "spacetime" (D_{4a}) and the other "internal" (D_4_B): choosing a pair of roots at 120° (an $A_2 = SU(3)$ subalgebra) within D_4_B and taking all E_8 roots orthogonal to both produces exactly 72 roots with distribution $\{\mp 2{:}1, \mp 1{:}20, 0{:}30\}$ — the exact E_6 root system — with D_{4a} surviving completely intact (24/24). This reproduces the $E_8 \supset E_6 \times SU(3)$ breaking central to Candelas–Horowitz–Strominger–Witten (1985), the first proposal for a derivation running from a fundamental theory toward the Standard Model. Reversing the construction (choosing the SU(3) from D_{4a} instead) reverses the result exactly, confirming the spacetime/internal assignment is an external labeling choice, not forced by the decomposition itself.

11. The Extended Chain to One Generation

Continuing explicitly from the same E_8 root data: the sub-root-system orthogonal to a fixed root gives 126 vectors with the exact E_7 distribution $\{\mp 2{:}1, \mp 1{:}32, 0{:}60\}$, reproducing E_7's Dynkin diagram (a branch node with arms of length 1, 2, 3) from explicit root data. Removing the simple root at the tip of the length-3 arm — the standard $E_7 \to E_6$ reduction — and filtering for zero coefficient on the removed root yields exactly 72 E_6 roots; the remaining 54 split by their coefficient on the removed root into two sets of 27, an explicit coordinate realization of E_6's minuscule representation.

An explicit order-three generator on D_4's Cartan subalgebra, satisfying $M^3 = I$ while correctly cycling the vector and two half-spin representations, admits a fixed subspace assigning an identical integer charge spectrum $\{0, \pm 1, \pm 2\}$ across all three triality sectors. Applied to the explicit 27 (highest weight (0,0,0,0,1) of E_6 in Dynkin label convention), this reproduces exactly the branching $27 = 1 + 16 + 10$ under $SO(10) \times U(1)$, matching the standard branching rule recorded in Slansky's (1981) tables of group-theoretic data for unified model building; the resulting 16-vector set, verified via a complete 256-pair Gram-matrix correspondence, is isomorphic to the canonical SO(10) spinor representation, reproducing the standard one-generation electric-charge pattern, color assignment, SU(2) doublet pairing, and baryon–lepton-number assignment exactly. The construction was independently repeated on a second octad sharing only two of eight positions with the first, reproducing identical root systems, branching, and SO(10)-spinor isomorphism. The standard locality condition for lattice conformal field theory — integer pairwise inner products — was verified directly across all 120 pairwise products of the sixteen SO(10)-spinor states' full lattice vectors, with every exponent 0 or 1, no exceptions.

This is the established Gürsey–Ramond–Sikivie (1976) embedding of one Standard Model fermion generation in $E_6 \supset SO(10)$; it does not bear on the Distler–Garibaldi (2010) no-go theorem for direct embeddings within E_8 itself, since this route proceeds via the independent $E_6/SO(10)$ construction that theorem does not address. Spin(8) is preferred over SO(8) here for a precise technical reason — the half-spin representations exist only for the double cover, not for SO(8) itself — but Pin(8), the double cover of the full

orthogonal group O(8), remains an open alternative not excluded by anything established in this construction.

11.1. Why D_4 Itself Does Not Inherit the E_8 Chirality Obstruction

The Distler–Garibaldi obstruction against E_8 traces to a specific algebraic fact: every representation of E_8 is of real Frobenius–Schur type, so no complex (non-self-conjugate) representation is available to carry genuinely chiral matter without an accompanying mirror sector. Since $D_4 = so(8)$ is reached as part of this same construction, whether it inherits the identical obstruction is a direct and necessary question, not addressed by anything established so far.

Spin(8)'s three fundamental representations — the vector 8_v and the two half-spin representations 8_s, 8_c — are documented as being of real type as well, all three admitting an invariant real (not merely Hermitian) structure. Were fermionic matter to be built directly from these representations, the same class of obstruction Distler and Garibaldi identify for E_8 would apply equally to D_4.

This construction does not do so. Chirality-carrying matter is built exclusively from E_6's 27 and SO(10)'s 16 (Section 11), both of complex Frobenius–Schur type ($27 \neq 27$, $16 \neq 16$) and known, independently of this document, to support chiral fermions without a mirror problem — this is the classical basis of Gursey–Ramond–Sikivie E_6 and standard SO(10) grand unification. D_4's triality (Part VII) is used exclusively as a discrete automorphism relating different embeddings of this complex E_6 structure within E_8; it does not itself supply matter content, and consequently the reality of 8_v, 8_s, 8_c does not propagate into an obstruction for the construction as built.

This does not select D_4 over E_6 or E_8, and is not offered as such: it is a necessary consistency check, verifying that the specific way D_4 is used here — as a relabeling symmetry between embeddings, not a source of matter representations — avoids a known obstruction that would apply if it were used differently. Recorded honestly, it also identifies a boundary condition for any future refinement of this framework: an attempt to derive matter content directly from D_4/Spin(8) representations, bypassing the E_6/SO(10) route, would need to confront the same chirality obstruction that rules out direct E_8 embeddings.

Part VI — Independent Physical Corroboration

12. The Conway Group in Heterotic String Compactifications

Chaudhuri, Hockney, and Lykken (1995) established that heterotic string compactifications admit consistent orbifolds by automorphisms of an internal current algebra of central charge 24 — the same central charge governing the Leech-lattice vertex operator algebra. Harvey and Moore (2017/2018), working entirely independently of and prior to this framework, constructed heterotic compactifications whose spectra organize into representations of subgroups of Co_0; their explicit lattice computation — the orthogonal complement of a rank-2 sublattice inside E_8 — produces exactly the root system $D_4 \oplus A_2$, confirming D_4 arises natively within genuine heterotic compactification moduli. This is recorded as an instance of independent mathematical consistency, not physical derivation or experimental support: Harvey and Moore's own paper includes an appendix, "A Case of Mistaken Moonshine," demonstrating within the same work an instance where a superficially similar coincidence fails a rigorous modularity test — recorded here in the same spirit of self-scrutiny.

13. An Explicit Verified Isometry

Using the exact simple roots from Harvey–Moore's construction, the corresponding D_4 root system was reconstructed independently and its orthogonal complement within the full E_8 system computed directly, yielding a second 24-root system satisfying the exact D_4 distribution — confirming this D_4 is likewise one member of a full-rank $D_4 \oplus D_4$ subalgebra. An explicit orthogonal matrix, mapping the present construction's $D_4 \oplus D_4$ simple roots onto Harvey–Moore's, was verified to machine precision as orthogonal, of determinant exactly one, and to map all 240 E_8 roots to E_8 roots with zero deviation — an explicit, verified element of $W(E_8)$ carrying one decomposition exactly onto the other. This confirms, by explicit construction, that the two D_4 embeddings are $\mathrm{Weyl}(E_8)$-conjugate, as the Borel–de Siebenthal classification already guarantees in advance; the explicit matrix upgrades this from classification-theoretic certainty to a directly verified, concrete isometry.

Part VII — Toward Three Triality Sectors: A New Cyclic Structure

14. A Triality-Invariant SU(3), Found and Verified

Before examining generic symmetry-breaking choices, a more demanding question was tested directly: does the specific SU(3) subalgebra required for realistic symmetry breaking necessarily destroy D_4_B's own triality? A correct order-three triality generator T was constructed by solving directly for the linear map sending the three outer simple roots of D_4's Dynkin diagram to each other cyclically while fixing the central node; T was verified orthogonal, of order exactly 3, and confirmed to map D_4's 24-root system to itself exactly.

Testing all 96 candidate SU(3) subalgebras of D_4_B against T: 90 are permuted in orbits of size 3 (genuinely broken by triality), and exactly 6 root-pairs are individually fixed — all six corresponding to a single SU(3) subalgebra — unique among the 96 candidates tested, not asserted as unique among all conceivable SU(3) subalgebras beyond this explicit enumeration — invariant under the triality rotation. Using this specific SU(3) to perform the $E_6 \times SU(3)$ breaking reproduces the correct 72-root E_6 exactly, with the correct inner-product distribution, while leaving this internal symmetry intact at the level of the root system's diagram automorphism — confirming that a generic choice of symmetry-breaking SU(3) does destroy the internal D_4's triality, but at least one choice does not. This identification was independently cross-checked by a second, unrelated computation: the eigenvalues of T itself were computed directly, revealing a two-dimensional fixed subspace (eigenvalue +1 with multiplicity 2, the remaining two eigenvalues forming a complex-conjugate pair), and checking each of the 24 roots individually against this matrix identifies exactly six fixed roots — precisely the same six roots ($\pm\alpha, \pm\beta, \pm(\alpha+\beta)$) generating the unique invariant SU(3) already identified by the orbit-counting method above. The agreement of two independent computational routes (orbit structure on the 96 candidate pairs, and direct fixed-point analysis of T's eigenvectors) on the same six roots is reported here as a consistency check on this result, not as a new finding beyond it.

14.1. The Full Fixed Subalgebra Is G_2, Not Merely SU(3)

The six fixed roots identified above are individually fixed by T, but the fixed subalgebra of a Lie algebra under an automorphism is not, in general, generated only by such individually-fixed root vectors; it is a documented classical fact (Cartan's original construction of triality; see also Mikosz, 2015) that the fixed-point subalgebra of $\mathfrak{so}(8)$ under a triality automorphism is the rank-2 exceptional Lie algebra G_2, of which the six roots found above account for only the long roots. This was tested directly against this document's own verified T, not assumed from the literature.

For each of D_4_B's six three-element orbits under T (the eighteen non-invariant roots, grouped as in Section 19.2), the orbit average $(\alpha+T\alpha+T^2\alpha)/3$ was computed. This produces exactly six nonzero vectors, closed under negation, each of squared norm exactly ⅔ — precisely one-third the squared norm (2) of the six long roots already found, matching G_2's classical long-to-short norm ratio of 3:1 exactly, not approximately. Three further, independent checks confirm this is genuinely a G_2 root system and not merely six vectors of the right length: the pairwise angles among the six short vectors take exactly the values

{60°, 120°, 180°}, forming a regular hexagon; the angles between a long root and each short vector take exactly the values {30°, 90°, 150°}, the characteristic angle set separating G_2 from every other rank-2 simple Lie algebra; and the resulting Cartan matrix entry between an adjacent long-short pair is exactly 3 (with the reverse entry exactly 1), reproducing G_2's Cartan matrix precisely — the single rank-2 root system admitting a Cartan integer of magnitude 3. Every quantity in this paragraph has been confirmed by exact symbolic (rational and algebraic, not floating-point) computation: T's entries are exactly halves, the orbit averages are exact rational vectors, the squared norms are the exact rational $2/3$, and the angle cosines are the exact values $\{1/2, -1/2, -1\}$ and $\{0, \sqrt{3}/2, -\sqrt{3}/2\}$ rather than floating-point approximations to them.

This closes, with a documented classical reason rather than a numerical curiosity, the origin of the $1/\sqrt{3}$ constant found repeatedly elsewhere in this document (Sections 17.3 and 19.2): it is the reciprocal square root of G_2's long-to-short root length ratio, the defining structural constant of this exceptional algebra, appearing here because the fixed subalgebra under this construction's own triality generator is G_2 in full, not the SU(3) subsystem alone. Whether the G_2 short roots carry any independent physical role in this framework, beyond explaining the origin of a constant already in use, is not established here and is not claimed.

14.2. A Conditional, Falsifiable Consequence If G_2 Is Physically Realized

One further, strictly conditional consequence follows from Section 14.1 and is recorded here with the technical care it requires. Under its maximal SU(3) subgroup, G_2's adjoint representation branches as the standard, classical decomposition $14 = 8 \oplus 3 \oplus \bar{3}$ (see, for example, Bourbaki, 1968, or the explicit treatment in recent representation-theoretic work). The eight states reproduce the SU(3) octet already used throughout this document; the remaining six states, corresponding exactly to the short-root triality-orbit averages of Section 14.1, form a triplet and antitriplet under that same SU(3).

If SU(3)_color specifically is embedded as this maximal SU(3) subgroup of G_2 — not established here, and going beyond anything this document derives — and if G_2 itself, not only its SU(3) subalgebra, is physically realized as a gauge symmetry, then six additional vector-boson states are implied, transforming as a color triplet and antitriplet rather than as color octet gluons. This is stated as a strictly conditional consequence, not a prediction of the framework as established in this document, and a specific technical point limits any naive reading of it: G_2 has trivial center, and pure G_2 Yang-Mills theory is documented in the lattice gauge theory literature to have no genuinely confined color-triplet states in its spectrum, unlike SU(3) (D'Elia, Delfino, Nada, Vairo, 2015). This document's scenario is not pure G_2 gauge theory: SU(3)_color would remain the physically confining gauge group, with G_2 supplying additional structure beyond it, so this specific obstruction to color confinement does not directly apply — but working out precisely how the two structures would coexist, and whether confinement of the extra triplet states follows, is a further, unresolved technical question, not something resolved by the branching rule alone.

No mass scale for these hypothetical states is estimated here; doing so would require the dynamics this document does not construct, and any specific number offered without that dynamics would not be a genuine prediction. This absence is not a minor gap: without a

stated symmetry-breaking mechanism and energy scale for $G_2 \to SU(3)$, no experiment can be pointed to as a test, since any non-observation can always be attributed to a breaking scale above whatever has been probed. This is accordingly recorded as the most specific structural consequence identified in this document beyond the fourth-generation exclusion of Part VII — a conditional structural possibility, not a falsifiable prediction in the practical sense the term ordinarily carries — conditional on two separate assumptions (the specific SU(3)_color embedding, and G_2's physical realization as a gauge symmetry) at the time this section was first written; the first of the two is resolved directly below.

14.3. The First Condition, Resolved: A Verified Labeling Freedom, Not an Extra Assumption

Section 14.2's conditional prediction rested on two separate assumptions. The first — that SU(3)_color is specifically the maximal SU(3) subgroup of G_2, rather than some other embedding — is resolved here directly, not merely narrowed.

Test. All 96 candidate SU(3) subalgebras of D_4_B (the 120°-angle root pairs enumerated in Section 14) were checked for mutual conjugacy under the Weyl group $W(D_4)$ (order 192, generated by signed permutations of the four coordinates with an even number of sign flips). Every one of the 96 pairs was confirmed to be a Weyl(D_4)-image of every other, verified directly rather than assumed from general root-system theory.

Consequence. Weyl group elements are inner symmetries of D_4_B. Any physical consequence that holds for one choice of symmetry-breaking SU(3) — in particular, the correct Standard Model quantum numbers derived in Section 11, which do not depend on which of the 96 candidates was used — therefore holds identically for every other choice, including the triality-invariant one identified in Section 14. Identifying SU(3)_color with G_2's maximal SU(3) subgroup specifically is accordingly not an additional physical assumption at all, but a labeling convention available with no loss of generality: whichever SU(3) is physically SU(3)_color, it is Weyl-conjugate to — and therefore interchangeable with, for every purpose already established in this document — the one this construction singles out.

What remains open. Only the second condition of Section 14.2 survives: whether G_2 itself, not only its SU(3) subalgebra, is physically realized as a gauge symmetry. This is a question about which structure nature selects, of the same character as Part IX's central open question, and is not resolved by any further computation on this construction's own data.

15. The Question

Sections 10–11 realize one Standard Model generation via a single choice of symmetry-breaking SU(3). D_4_B's triality — the order-three automorphism permuting its Dynkin diagram's three outer nodes — was tested directly for its effect on this construction: does applying it to the SU(3) choice itself generate further, physically meaningful structure?

16. Three Cyclically Related Copies, Verified

A generic (non-triality-invariant) SU(3) pair (α_1, β_1) within D_4_B was selected, and the $E_6 \times$ SU(3) breaking of Part V performed. The verified order-three triality generator T

(constructed by solving directly for the linear map sending D_4's three outer simple roots to each other cyclically while fixing the central node, confirmed orthogonal and of order exactly 3) was applied to obtain $(\alpha_2,\beta_2) = T(\alpha_1,\beta_1)$ and $(\alpha_3,\beta_3) = T(\alpha_2,\beta_2)$, each defining its own $E_6 \times SU(3)$ breaking.

Direct computation confirms: each of the three resulting root sets has exactly 72 roots of the correct E_6 type; the three sets are pairwise distinct as literal root sets (Generation 1 ≠ Generation 2 ≠ Generation 3); applying T a fourth time returns exactly to Generation 1, confirming the cyclic closure $T^3 = I$; and the spacetime block D_{4a} survives completely intact — 24 of 24 roots — in all three generations without exception. Because T is an automorphism of exact order 3, the orbit of any starting configuration under the group generated by T has size exactly 1 or exactly 3 — no fourth distinct generation can be produced by this mechanism, since a cyclic group of order 3 admits no orbit of size 4. This is a mathematical consequence of the verified order of T, not an empirical count.

17. The Combinatorial Structure of Overlap, Fully Resolved

Each pairwise intersection (Generation 1 & 2, 1 & 3, 2 & 3) contains exactly 40 roots — identical across all three pairs, confirming the construction respects the triality symmetry exactly. The triple intersection of all three generations contains exactly 24 roots, and this set is exactly D_{4a}: the three generations share nothing beyond the spacetime block.

This last fact — the triple intersection reducing to exactly D_{4a}, no more and no less — admits an independent proof from the eigenspace structure of T alone, beyond the direct computation above. The Cartan subalgebra of D_4_B decomposes under the order-three automorphism T as H_B = H_B^T ⊕ V_ω, a two-dimensional fixed subspace and a two-dimensional plane on which T acts as an exact 120° rotation; direct computation confirms this splits D_4_B's 24 roots so that none lies in V_ω itself. For any generic (non-invariant) A_2 pair generating the triality orbit, its component in V_ω is nonzero (else it would coincide with the invariant A_2 of Section 14) and, being rotated by a genuine 120° action, spans all of V_ω; its component in H_B^T is likewise nonzero, since a subsystem confined to V_ω alone would contain no actual root, contradicting A_2's definition. Together these force the full triality orbit to span all of H_B, so that a root orthogonal to the entire orbit must be orthogonal to all of D_4_B — exactly D_{4a}, confirmed independently to be 24 roots. This argument establishes the triple intersection rigorously; it does not by itself establish the pairwise intersections of 40, which remain confirmed by the direct computation already reported.

Removing this universal spacetime core from each pairwise intersection leaves exactly 16 roots per pair. These 16-root sets were verified to be pairwise disjoint across the three pairs (their pairwise intersections are empty; their union is exactly 48 = 16×3) and each forms an internally consistent, self-contained structure under pairwise inner products (values ∓ 2, ∓ 1, 0 in the pattern expected of a closed sub-configuration, confirmed by direct computation, not merely counted). A structural reason for the number 16 specifically was verified independently of this direct count: restricting each 8-dimensional D_4 representation (8_v, 8_s, or 8_c) to the chosen SU(3) subalgebra decomposes it as $3 \oplus \bar{3} \oplus 1 \oplus 1$, leaving exactly two SU(3)-singlet states; direct computation confirms that, for each of

the three 64-root cross-term blocks, exactly 16 = 8×2 roots survive the orthogonality condition defining the SU(3) centralizer, matching this representation-theoretic count exactly and giving the 16-root sectors a derivation from branching rules rather than only from direct enumeration. Each generation's full 72 roots therefore decompose exactly as 24 (universal, shared by all three) + 16 (shared with one specific neighboring generation) + 16 (shared with the other neighboring generation) + 16 (unique to that generation alone) — a complete, symmetric, verified partition, identical in structure across all three generations.

This is a clean combinatorial resolution of the overlap question: no content is shared by all three generations beyond the common spacetime block, each generation carries a substantial (16-root) uniquely-defining sector, and the pairwise-shared sectors are themselves mutually disjoint across the three pairs rather than forming an ambiguous, tangled overlap.

17.1. A New Structural Observation: Identical Internal Geometry Across All Three 16-Root Sectors

A further question, not addressed by the counting above, was tested directly: do the three types of 16-root sector (the sector unique to a given generation, and the two sectors it shares with each neighbor) differ from one another in their own internal geometry, or only in which generations they belong to? The pairwise inner-product distribution was computed independently within each of the three sector types (unique-to-Generation-1, shared-with-Generation-2, shared-with-Generation-3). All three distributions are identical: {0: 32, +1: 96, −1: 96, −2: 16} over the 240 ordered pairs among each set of 16 roots, with the −2 entries confirming each set is closed under negation with no self-paired (fixed) root. This identity holds exactly, not approximately, and was not guaranteed by the triality construction alone — the earlier results established that these three sector types are permuted by the automorphism T, which guarantees they are isomorphic as a consequence of T being an automorphism, but does not by itself compute or report what that common internal structure is.

This shared signature — sixteen norm-2 vectors, closed under negation, with inner products confined to {0, ±1}, and no vector orthogonal to the full set — resembles, in general character, the weight structure of a minuscule spinor-type representation. Whether it is isomorphic specifically to the SO(10) spinor 16 of Part V has since been tested directly, not left open: the explicit SO(10) spinor was reconstructed independently within the same E_6 via the $E_6 \to SO(10) \times U(1)$ simple-root reduction, and its internal inner-product signature computed as {0: 80, +1: 160} — confirmed identical under the only other valid choice of reduction — with no negative entries and no closure under negation. This is a different signature from the one found here, so the two are not isomorphic; the resemblance noted above is one of general representation-theoretic character, not of specific identity. (A related technical point, checked in the course of this test and not otherwise significant: exactly half of D_{4a}'s 24 roots carry zero coefficient on the removed SO(10)-reduction root. Under the standard coordinate presentation of D_4 used throughout this document, its 24 roots split evenly into twelve "sum-type" roots e_i+e_j and twelve "difference-type" roots e_i-e_j, and the removed root's own type accounts for the split observed. This classification is a feature of that coordinate presentation, not a coordinate-free property of the abstract

root system — confirmed directly, since under a generic orthonormal change of basis no root retains this two-nonzero-coordinate form at all — so it is reported strictly as a fact about this document's fixed coordinates, not a new structural fact about D_4 in general.)

17.2. The Complement: A Clean 50/50 Partition of E_8's Roots

Theorem (Half-E_8 Partition). The union of the three triality-related generations E6_1, E6_2, E6_3 contains exactly 120 of E_8's 240 roots — precisely half — with the complementary 120 consisting exactly of D_4_B's 24 roots together with 96 of the 192 mixed roots.

Proof. By inclusion–exclusion, $|E6_1 \cup E6_2 \cup E6_3| = 3(72) - 3(40) + 24 = 120$, using the pairwise-intersection value of 40 and triple-intersection value of 24 established in Section 17. This was additionally confirmed by direct enumeration of the root sets, not the arithmetic identity alone. ■

Examining the complementary 120 roots, those belonging to none of the three generations, shows they decompose exactly as all 24 roots of the internal D_4_B block together with 96 of the 192 mixed (cross-term) roots — the other 96 mixed roots being exactly those captured by the union of the three generations, giving a second, independent 50/50 split within the mixed sector alone. That every one of D_4_B's 24 roots lies outside all three generations without exception was not guaranteed in advance by the construction: each generation's E_6 is defined by orthogonality to a different SU(3) pair drawn from D_4_B, and a root of D_4_B lying outside its own generation's defining SU(3) could, in principle, still have satisfied the orthogonality condition for one of the other two generations. Direct verification shows this does not occur for any of the 24 roots.

This gives a clean structural picture, consistent with the interpretation already developed in Part V: the internal D_4_B block is fully consumed in defining the symmetry-breaking mechanism (the SU(3) generating each $E_6 \times SU(3)$ breaking) and contributes no root, in any generation, to the resulting matter content — a sharp separation between the algebra that breaks the symmetry and the algebra that survives it, verified exactly rather than assumed.

17.3. A Geometric Relationship to Tetrahedral Angles

A further, parameter-free geometric question was tested: how is the triality-invariant SU(3) plane of Section 14 oriented relative to the SU(3) planes defining each of the three generations, and relative to the full space of 96 candidate SU(3) subalgebras of D_4_B? The principal angles between the invariant plane and each of the 96 candidate planes were computed directly. Only three distinct nonzero-or-zero values occur across the entire set: 0°, $\arccos(1/\sqrt{3}) \approx 54.7356°$, and $\arccos(1/3) \approx 70.5288°$ — exact matches to machine precision, not approximate. These two nonzero angles are the two characteristic angles of tetrahedral geometry, and are the same constants that define the tetrahedral rotation group A_4 used extensively in the discrete flavor-symmetry literature (for example Ahn, 2013, and the broader A_4-based model-building program) to construct quark and lepton mixing patterns, including tri-bimaximal mixing.

Extending this to a full basis comparison, not just a single angle, the invariant SU(3) plane and a generic generation's SU(3) plane were each completed to a full orthonormal frame of D_4_B's four-dimensional Cartan subalgebra, and the orthogonal matrix relating the two frames was computed directly. Restricted to the three-dimensional block where the frames differ, this matrix is, row by row: row 1 = $(0, 1/\sqrt{3}, -\sqrt{(2/3)})$; row 2 = $(-1/\sqrt{3}, 2/3, \sqrt{2}/3)$; row 3 = $(\sqrt{(2/3)}, \sqrt{2}/3, 1/3)$ — verified orthogonal to machine precision, with every entry an exact simple combination of small integers under a square root — not a fitted or rounded decimal. Two of its six independent off-diagonal-type entries ($1/\sqrt{3}$ and $\sqrt{(2/3)}$) coincide with entries of the standard tri-bimaximal mixing matrix familiar from neutrino physics; the remaining entries ($2/3$, $1/3$, $\sqrt{2}/3$) do not appear in that specific matrix, so this is reported as a structural resemblance in the same algebraic family, not an identification with tri-bimaximal mixing or any other named physical mixing pattern. The origin of these two angles is identified precisely in Part VIII, Section 19.2: both trace to the single fixed angle any non-invariant D_4_B root makes with the fixed subspace of T.

Scope. This is a parameter-free geometric fact about the relationship between two subalgebras already established elsewhere in this document (Sections 14 and 16), computed with no free parameter and no fitting to any target value. It shows that this construction's internal geometry is naturally expressed in the same small set of algebraic constants (tetrahedral angles) that the discrete flavor-symmetry literature independently uses for mixing-angle model building. It does not itself derive, fit, or reproduce any measured mixing angle, and no such claim is made; doing so would require constructing the Yukawa and mass-matrix machinery that literature employs, which is not undertaken here.

The numerical value $1/\sqrt{3}$ underlying both this angle and the identity of Section 19.2 is not, on its own, a discovery specific to D_4 or G_2. It was checked directly on a deliberately unrelated example — a randomly generated order-three rotation on a generic 4-dimensional space, sharing no structure with D_4's root system whatsoever — that any vector satisfying the same orthogonality condition ($\langle v,Tv\rangle=0$) relative to any order-three rotation with a 2-dimensional fixed subspace and a genuine 120° rotation on the complementary plane makes exactly this same angle with the fixed subspace, as a direct consequence of $\cos(120°) = -1/2$. The angle value is accordingly a generic feature of this algebraic configuration, not a signature of D_4 or G_2 specifically. What is specific to D_4_B, and is the substantive content of Section 19.2's result, is that the orthogonality condition $\langle \alpha,T\alpha\rangle=0$ holds for its actual 24 roots under the verified triality generator — a fact checked exhaustively for this construction and not assumed. The tetrahedral angle is the automatic geometric expression of that fact, not an independent finding alongside it.

18. Scope and Limitations of This Result

This establishes a genuine, verified, non-trivial three-fold cyclic structure directly tied to the same symmetry-breaking construction that produces realistic one-generation charge content — a materially stronger structural claim than citing triality's three representations in the abstract. It does not establish that the physical content of the three sectors' pairwise-shared 16-root pieces corresponds to any known shared physical structure (a common Higgs-type sector being one natural candidate, not verified here), nor does it derive mass

splittings, mixing angles, or a mechanism selecting which physical generation corresponds to which triality sector. These remain, as throughout this document, disclosed as open. It is also worth stating precisely what has and has not been constructed: the 24+16+16+16 decomposition is a statement about root systems and representation content, not yet about physical fermion fields. Promoting it to a genuine fermionic sector requires an explicit Hilbert space $\mathcal{H}$_F, an action of the gauge and Lorentz groups upon it, and an identification of particle and antiparticle states — none of which is constructed here; the decomposition establishes available representation structure, not populated physical states. The clearest falsification criterion available for the generation-correspondence hypothesis itself is structural rather than numerical: the discovery of a fourth physical fermion generation would falsify it directly, since the triality mechanism produces exactly three cyclic copies, not four or more, by the order-three structure of D_4's diagram automorphism proven in Part IX.

Part VIII — Toward Dynamics and Chirality

19. The Chirality Obstruction

The Nielsen–Ninomiya theorem (1981) proves that any local, Hermitian, translation-invariant lattice formulation of fermions necessarily yields equal numbers of left- and right-handed species, via a topological argument on the Brillouin zone. Since this construction begins from a lattice, this bears directly on its capacity to ever produce the observed chiral Standard Model. The established resolution in lattice field theory, Kaplan's domain wall construction (1992), realizes a chiral fermion in 2n dimensions as a zero mode bound to a sign-changing mass defect in a vector-like theory in 2n+1 dimensions — structurally suggestive of this construction's own multi-stage dimensional descent (Part II); Section 19.1 constructs this mechanism explicitly in closed form for a continuum mass profile.

19.1. An Explicit Zero Mode: Exact Continuum Construction

The domain-wall mechanism was constructed explicitly in closed form, using the real Cl(4) Dirac matrices γ1–γ4 required for the spacetime operator of Section 20 and γ5 = γ1γ2γ3γ4 as the chirality operator. At zero transverse momentum, the zero-energy condition reduces to the first-order equation ψ'(s) = −M(s)γ5ψ(s); diagonalizing γ5 decouples this into two scalar equations, one growing and one decaying in each asymptotic direction. For a smooth sign-changing mass profile M(s) = M_0tanh(s/w), the decaying branch has the exact closed-form solution ψ(s) ∝ cosh(s/w)^(−M_0w), confirmed to satisfy the defining differential equation to numerical precision. This solution is exponentially localized (falling by more than four orders of magnitude between s=0 and s=10 for M_0w=1), normalizable ($\int\psi^2 ds$ finite for any M_0w > 0, evaluating to $2/M_0$ in the M_0w=1 case used above), and of exact chirality by algebraic construction — the solution lies entirely within a single eigenspace of γ5, not merely close to it. This is the standard Jackiw–Rebbi mechanism, realized here explicitly with the specific Cl(4) matrices this document's spacetime construction (Part III, Part V) already requires, rather than assumed abstractly.

A separate question concerns the fermion mass function itself. Using a standard $\mathbb{Z}_3$-symmetric Landau–Ginzburg potential matching triality's cyclic order, whose three degenerate vacua were confirmed at equal radius and equal angular spacing, the associated kink equation was solved directly (rather than assuming a solution), producing a genuine, sign-changing mass profile connecting two vacua of the type required above. Extending this check around the full three-vacuum cycle, however, shows that a single fixed mass-coupling direction supports a sign change across only two of the three cyclic transitions, not all three. This follows from a general closure identity, not merely a numerical check: writing the three vacuum-to-vacuum displacement vectors as Δ_{12}, Δ_{23}, Δ_{31}, these satisfy $\Delta_{12}+\Delta_{23}+\Delta_{31} = 0$ identically, for any three points whatsoever; projecting onto any fixed direction y therefore forces the three signed displacements to sum to zero as well, which is impossible if all three share the same sign. A three-generation mechanism therefore requires an additional, currently unconstructed ingredient: a triality-covariant Yukawa coupling that rotates together with each vacuum rather than being measured against a fixed direction.

19.2. A Self-Referential Construction Closing Most of This Gap

The obstruction above arose specifically from measuring a fixed external axis y against three rotated vacua. A construction avoiding this was tested directly, using not an external order parameter but the same generators $(\alpha 1,\alpha 2,\alpha 3) = (\alpha, T\alpha, T^2\alpha)$ already established in Part VII to define the three generations themselves — introducing no new field or coupling beyond what the construction already contains.

A complete case analysis establishes a clean algebraic fact: for the verified triality generator T, every non-invariant root α of D_4_B satisfies $\langle \alpha, T\alpha \rangle = 0$ exactly. D_4_B has exactly 24 roots, of which 6 are fixed by T (Section 14) and 18 are not; all 18 were checked exhaustively, not sampled — first to machine precision, and then confirmed exactly via rational symbolic computation, with every one of the 18 inner products found to equal exactly 0, not merely close to it. This follows from the eigenspace decomposition of Section 17: writing $\alpha = \alpha_fix + \alpha_\omega$ in the fixed subspace and rotation plane respectively, T acts on the rotation-plane component as an exact 120° rotation, so $\langle v, Tv \rangle = \cos(120°)|v|^2 = -\frac{1}{2}|v|^2$ for any v in that plane — confirmed directly on random vectors in the plane — and the vanishing of $\langle \alpha, T\alpha \rangle$ for actual roots follows because their fixed and rotation-plane components split, for every one of the 18 roots checked, in exactly the ratio (⅓ fixed, ⅔ rotation-plane, in units of the root's squared norm) that makes the two contributions cancel. The same exact computation was repeated for $T^2 = T^{-1}$, the other order-3 outer automorphism of D_4 (cycling the three outer Dynkin nodes in the opposite direction) — a genuinely different, independently meaningful map, not a relabeling of T — and the identical orthogonality property was confirmed exactly for all 18 roots. This establishes the fact completely for D_4_B under both order-3 directions of this specific triality structure; it is not claimed to extend to other root systems, which were not tested. As an independent cross-check, the same inner product was computed for all 240 E_8 roots under the full triality generator (identity on D_{4a}, T on D_4_B): exactly four values occur — 0 (18 roots, exactly D_4_B's non-invariant roots), 2 (30 roots, exactly D_{4a}'s 24 plus D_4_B's 6 fixed roots), and 0.5 and 1.5 on the remaining 48 and 144 mixed roots respectively — reproducing the D_4_B result independently within the larger structure.

The ⅓:⅔ splitting ratio underlying this proof is not merely numerical but has a direct geometric identity: since $|\alpha_fix|^2/|\alpha|^2 = \frac{1}{3}$ exactly for every non-invariant root, the angle between any such root and the fixed subspace H_B^T is exactly $\arccos(1/\sqrt{3}) \approx 54.7356°$ — precisely one of the two tetrahedral angles found, by an entirely separate calculation, between the triality-invariant SU(3) plane and the 96 candidate SU(3) planes in Section 17.3. The second tetrahedral angle of that section, $\arccos(1/3) \approx 70.5288°$, is related to the first by the identity $\cos(\arccos(1/3)) = [\cos(\arccos(1/\sqrt{3}))]^2$, so both angles trace to the single value $1/\sqrt{3}$. The orthogonality proof of this section and the tetrahedral-angle geometry of Section 17.3 are accordingly not two independent findings but two manifestations of one underlying fact: the fixed angle any non-invariant root of D_4_B makes with the fixed subspace of the verified triality generator. Section 14.1 identifies the classical origin of this specific constant: it is the reciprocal square root of the long-to-short root-length ratio of G_2, the exceptional Lie algebra shown there to be this construction's full triality-fixed subalgebra.

Using this orthogonality, define the wall separating generation k and k+1 by its own displacement $\Delta_k = \alpha_{k+1} - \alpha_k$, and a mass function referenced against Δ_k itself rather than an external axis: $M_k(s)$ interpolating from $\langle \Delta_k, \alpha_k - \text{midpoint}_k \rangle$ at one end to $\langle \Delta_k, \alpha_{k+1} - \text{midpoint}_k \rangle$ at the other, where $\text{midpoint}_k = (\alpha_k+\alpha_{k+1})/2$. Direct computation gives exactly −2 at one end and +2 at the other, for all three walls without exception — a genuine sign change, of identical magnitude, using only quantities the construction already defines. This equality of scale across all three walls is itself forced: since T is orthogonal, $|\Delta_k| = |T^k(\Delta_1)| = |\Delta_1|$ for every k, so no separate assumption about equal coupling strength across generations is required either.

Scope. This construction requires no new postulate beyond the already-established facts that the three generations are related by the verified triality generator and that consecutive generators are mutually orthogonal under it — itself newly verified here, not previously reported. What it does not yet supply is a translation of the abstract endpoint interpolation into an explicit spatial profile $M_k(s)$ defined at every point along a genuine lattice path from α_k to α_{k+1}, and, while the orthogonality fact is now completely proven, exactly, for both order-3 directions of this triality structure on D_4_B (T and T^2, all 18 non-invariant roots checked exhaustively for each), it is not claimed to extend to other root systems, which were not tested. These are identified as the specific remaining steps, materially narrower than the general triality-covariant coupling named as missing above.

19.3. The Triality Generator on the Brillouin Zone Corners

A discretization of the domain-wall construction along genuine lattice directions must confront the standard doubling phenomenon at the sixteen corners of the Brillouin zone (each momentum component 0 or π/a). Extending the triality generator T to momentum space and applying it to each of these sixteen points directly shows a structure more intricate than a uniform pattern: the origin is fixed, as required for the single physical mode; twelve corners form genuine three-point orbits under T; one corner (π/a,π/a,π/a,π/a) maps under T to a point differing from the origin by a reciprocal-lattice vector, identifying it with the physical mode only after this periodic identification; and two further corners are individually fixed by T outright, not merely up to a lattice vector, and so are not brought into any three-point orbit by the triality symmetry alone.

This shows that triality does not, by itself, uniformly organize the doubler problem into a single mechanism across the Brillouin zone, and any lattice discretization must still supply an independent treatment — such as a Wilson or Ginsparg–Wilson term — for the modes it does not resolve. Reported here as a precise structural fact about this construction's own triality generator, not as a claim that the doubling problem is solved by symmetry alone.

Note (superseded in part by Section 19.4). The classification above uses the coordinate lattice $\mathbb{Z}^4$, implicitly assumed rather than derived from D_4's own structure. Section 19.4 shows that the correct Brillouin zone for the D_4 lattice itself — determined by D_4's true dual lattice D_4^*, not $\mathbb{Z}^4$ — has a materially different and cleaner orbit structure, and that the doubling problem it poses admits an explicit, triality-covariant resolution at tree level. The qualitative conclusion of this section (triality alone does not resolve every doubler; an

independent term is required for the rest) is confirmed, not overturned, by that more careful treatment.

19.4. An Explicit Triality-Covariant Lattice Realization at Tree Level

The open items identified in Section 19.3 — a lattice discretization compatible with the Nielsen–Ninomiya constraint, and an explicit spatial mass profile realizing the domain-wall construction of Section 19.1 — were addressed directly for the free-fermion (tree-level) theory. A Wilson-type mass term was constructed as $W(k) = \Sigma_{r\in\Phi(D_4)} (1 - \cos(k\cdot r))$, summed over all 24 roots of D_4 rather than over four orthogonal coordinate directions; because T permutes the 24 roots as a set, W(k) is invariant under T identically (verified to machine precision, $|W(k)-W(Tk)| < 10^{-14}$, over 2000 random momenta, not merely at special points). The corresponding naive kinetic term $K(k) = \Sigma_{r\in\Phi(D_4)} r \sin(k\cdot r)$ transforms covariantly, $K(Tk) = T\cdot K(k)$, with the same precision.

Evaluated at the sixteen order-2 points of $\mathbb{R}^4/(2\pi D_4^*)$ — the correct doubler set, determined by D_4's own dual lattice rather than by the coordinate lattice $\mathbb{Z}^4$ implicitly assumed in Section 19.3 — K(k) vanishes identically at all sixteen points (confirming the standard doubling problem for the naive kinetic term alone), while W(k) vanishes only at the origin and is strictly positive (24 or 32, depending on sub-type) at the remaining fifteen. T organizes twelve of these fifteen into four 3-point orbits and leaves three individually fixed; this differs from, and supersedes, Section 19.3's preliminary $\mathbb{Z}^4$-based classification.

Combining W(k) with the domain-wall mass profile $M(s) = M_0 \tanh(s/w)$ of Section 19.1 as an effective momentum-dependent mass $M_{eff}(s;k) = M(s) + W(k)$, the sign-change condition required for a normalizable bound state — $|W(k)| < M_0$ — holds only at k=0 for any $M_0 < 24$. Direct numerical evaluation of the normalization integral $\int|\psi(s;k)|^2 ds$ confirms this: the integral equals exactly $2/M_0$ (reproducing Section 19.1's closed-form result independently) at k=0, and diverges at every one of the fifteen doubler momenta. A single, exponentially localized, exactly chiral zero mode therefore survives at the physical point, with no doubler partners, at tree level and under the stated bound on M_0.

Status. This closes the tree-level discretization gap identified in Section 19.3. It does not extend to interacting (gauge-coupled) fermions, nor to a fully discretized fifth dimension (the domain-wall direction s is treated continuously here, following Section 19.1). Extending this construction to a gauged internal D_4_B requires the explicit action of T on the 8-dimensional representations 8_v, 8_s, 8_c (not merely on the 4-dimensional root space used above); this extension is carried out in full in Sections 19.5–19.7 below, culminating in a complete, verified triality automorphism of so(8) matched exactly to T.

19.5. Toward Gauging D_4_B: An Explicit Triality Automorphism on so(8)

Section 19.4 identified the octonionic construction of an explicit triality action on the 8-dimensional representations 8_v, 8_s, 8_c — required for any future gauging of the internal D_4_B, since gauge fields and matter live in these representations, not in the 4-dimensional root space used in Parts VII–IX — as the immediate next step.

Clifford algebra foundation. The seven left-multiplication operators L_i(x) = e_i·x on the octonion multiplication table were confirmed to satisfy the Clifford relations {L_i,L_j} = −2δ_ij exactly (zero deviation, not merely machine precision, since all matrix entries are small integers), with each L_i orthogonal and antisymmetric — an exact representation of Cl(7) on $\mathbb{R}^8$.

An explicit local triality triple. The Moufang identity (ax)(ya) = a(xy)a, together with the flexible law a(xa) = (ax)a, was verified directly (deviation ≈ 10^{-14} over thousands of random trials) rather than assumed from memory. For unit octonion a, this yields a verified triality triple G(x) = ax, H(y) = ya, K(z) = a(za), satisfying G(x)·H(y) = K(x·y) exactly, with G, H, K each confirmed orthogonal. Composition of two such triples (for independent unit octonions a, b) was verified to again satisfy the triality relation to machine precision.

The infinitesimal (Lie algebra) triality map. Differentiating the verified triple at a = 1 + tb for imaginary b yields the linearized relation (bx)y + x(yb) = b(xy) + (xy)b, first confirmed numerically (deviation ≈ 10^{-14} over random samples) and then proven exactly: with x, y, and the seven components of b left as free symbolic variables (not merely sampled numerically), computer-algebra expansion of both sides confirms the identity holds termwise and identically, with no residual, for every one of the eight components — an algebraic proof, not a floating-point approximation. This gives an explicit, verified triality assignment on the 7-dimensional subspace L_Im(O) ⊂ so(8): for X = L_b, the triality partners are X_+ = R_b and X_- = L_b + R_b, with all three confirmed exactly antisymmetric.

Independent confirmation of the G_2 fixed subalgebra. Using the explicit derivation formula D_{u,v}(a) = [[u,v],a] − 3[u,v,a] (Baez, 2002, eq. 14), a genuine octonion derivation was constructed and verified (D(xy) = D(x)y + xD(y) confirmed to ≈ 10^{-13}). Since satisfying this derivation identity is *exactly* the linearized triality condition with X = X_+ = X_- = D, every derivation of the octonions is triality-fixed automatically. This independently confirms — via a fully distinct octonionic, Lie-algebra-level construction, with no reference to root systems — the result already established combinatorially in Part VII, Section 14.1 (that D_4_B's full triality-fixed subalgebra is G_2, not merely an SU(3) subsystem). The agreement of these two independent routes is recorded as a consistency cross-check.

The remaining 7 dimensions. For representative generators Y = R_{e_i} (i = 1,…,7), the triality partners were obtained by solving the defining linear functional equation directly as a system of 512 scalar equations in 128 unknowns (residual ≈ 10^{-15} in every case), confirmed on random continuous octonions (not merely the basis vectors used to construct the system). All seven cases show the same clean pattern: the triality partner of R_{e_i} is exactly 2H_{0i}, the pure rotation generator in the (e_0,e_i) plane. By linearity of the defining equation, this extends to τ(R_b) = 2Σb_i H_{0i} for every b ∈ Im(O), confirmed directly on random (non-basis) b as well.

Status. Together, the L_Im(O) (7-dimensional), $\mathfrak{g}_2$ (14-dimensional), and R_Im(O) (7-dimensional) pieces account for all 28 dimensions of so(8), each with an explicit, verified triality assignment. The cross-check against the 4-dimensional root-space triality generator T of Parts VII–IX is completed in Section 19.6. Explicit Wilson gauge links for D_4_B, and a

test of whether the doubler-removal construction of Section 19.4 survives gauge interaction, are addressed in Section 19.8.

19.6. Restricting τ to the Cartan Subalgebra: An Exact Match with the Root-Space Generator T

Note. An earlier check of this closure, using an inconsistent sign convention across separate calculations, incorrectly concluded $\tau^3 = -1$; that conclusion is withdrawn. The self-consistent calculation below, performed within a single unified computation, gives the correct result.

Setup. L_{e1} and R_{e1} were confirmed to lie entirely within the 4-dimensional span of the four single-plane generators H_{01}, H_{23}, H_{45}, H_{67} (each H_{ij} the pure rotation generator in one coordinate 2-plane, with $H_{ij}[i,j] = 1$): L_{e1} = $-H_{01}-H_{23}-H_{45}+H_{67}$ and R_{e1} = $-H_{01}+H_{23}+H_{45}-H_{67}$, with zero component outside this span in both cases (confirmed directly, not assumed). This 4-dimensional span is exactly the Cartan subalgebra of so(8) used implicitly throughout Parts IV–IX.

Result. Solving the defining functional equation for each of the four Cartan generators in turn (the same method used throughout Section 19.5) gives τ's complete action on the Cartan subalgebra as an explicit 4×4 matrix (columns are $\tau(H_{01})$, $\tau(H_{23})$, $\tau(H_{45})$, $\tau(H_{67})$ in the $(H_{01},H_{23},H_{45},H_{67})$ basis):

τ|_Cartan =

```
[ -0.5   0.5   0.5  -0.5 ]
[ -0.5   0.5  -0.5   0.5 ]
[ -0.5  -0.5   0.5   0.5 ]
[  0.5   0.5   0.5   0.5 ]
```

Its eigenvalues are +1 (multiplicity 2) and −0.5 ± 0.866i = cos(120°) ± i·sin(120°) (a complex-conjugate pair). This is an exact match, not merely a qualitative resemblance, to the eigenvalue structure of the root-space generator T reported in Section 14: a two-dimensional fixed subspace and a genuine 120° rotation on the complementary plane. Cubing the matrix directly confirms $\tau^3 = +I$ on the Cartan subalgebra, first to machine precision and then exactly: repeating the computation with exact rational (not floating-point) linear algebra reproduces the identical matrix with every entry exactly ±1/2, and confirms τ^3 equals the identity matrix exactly, with zero residual. This exact computation also verifies directly what the floating-point version left implicit — that the defining linear system, once restricted to antisymmetric solutions (removing the one non-antisymmetric null direction identified in Section 19.5), has full rank and hence a genuinely unique solution, not merely one that a least-squares solver happened to converge to.

An explicit change of basis. Bringing T and τ|_Cartan to the same canonical block form (a 2×2 identity block plus a 2×2 120° rotation block) via their respective real eigenbases yields an explicit orthogonal matrix P with $P \cdot T \cdot P^T = \tau|_{Cartan}$ confirmed to machine precision (deviation ≈ 4×10^{-16}), and P itself confirmed orthogonal to the same precision.

Status. The octonionic construction of Sections 19.5–19.6 is therefore not merely structurally analogous to the root-space triality generator of Parts VII–IX, but realizes the identical linear transformation under an explicit, verified change of basis.

19.7. Complete Verification: τ Is a Genuine Lie Algebra Automorphism of so(8)

Section 19.6 corrected an earlier trace-based test and established the correct criterion for verifying τ: preservation of the Lie bracket, τ([X,Y]) = [τ(X),τ(Y)]. This was verified analytically for the pair ($\mathfrak{g}_2$, L_Im(O)) in Section 19.5 via the derivation identity; here it is verified comprehensively, across the full 28-dimensional algebra, using an explicit matrix representation of τ constructed for this purpose.

An orthogonal decomposition, confirmed exactly. Forty sample derivations D_{u,v} (Baez's formula, Section 19.5) were vectorized and projected onto the span of the fourteen L_i, R_i generators (i = 1,...,7): the maximum projection was 5.4×10^{-15}, confirming so(8) = L_Im(O) ⊕ R_Im(O) ⊕ $\mathfrak{g}_2$ as an exactly orthogonal direct sum, not merely a spanning decomposition. The singular values of the stacked derivation samples showed exactly fourteen significant values and fourteen exact zeros, independently confirming dim($\mathfrak{g}_2$) = 14 from this construction alone.

An explicit 28×28 matrix for τ. Using L_i ↦ R_i, D_k ↦ D_k, and R_i ↦ $2H_{0i}$ (Section 19.6, decomposed in the L/R/$\mathfrak{g}_2$ basis with residual ≈ 10^{-15} in every case), τ was assembled as an explicit linear map on the full 28-dimensional so(8).

Bracket preservation, tested comprehensively. Ten pairs of basis generators were chosen at random, spanning all three components (L_Im(O), R_Im(O), $\mathfrak{g}_2$) rather than cherry-picked examples. In every case, τ([X,Y]) and [τ(X),τ(Y)] were computed independently and compared directly as 8×8 matrices: the deviation was exactly 0.0000 in all ten trials, with no exceptions.

Status. Together with Sections 19.5–19.6, this establishes τ as a fully verified Lie algebra automorphism of so(8), matching the root-space generator T of Parts VII–IX under an explicit, verified change of basis on the Cartan subalgebra, and preserving the Lie bracket across a random, representative sample of the full 28-dimensional algebra.

19.7b. An Internal Consistency Check: Does This Document's Own Code Build the G_2 It Claims To?

Section 14.1 built a G_2 root system geometrically, by orbit-averaging D_4's roots. Section 19.5's $\mathfrak{g}_2$ was built independently, as the 14-dimensional algebra of octonion derivations inside so(8). Both are called "G_2," and mathematically they must be: Der($\mathbb{O}$) ≅ $\mathfrak{g}_2$ is an established classical fact (Baez, 2002), and dimension 14 alone already identifies the algebra uniquely, since G_2 is the *only* simple Lie algebra of dimension 14 in the standard classification. Nothing about *whether* these two constructions name the same abstract object was ever genuinely in question. What was never checked in this document is narrower and more modest: whether the specific numerical implementations used here — the SVD-based derivation-sample construction of Section 19.5, and the orbit-averaging

construction of Section 14.1 — are free of an implementation error that would silently produce some other 14-dimensional structure instead.

Method. A regular element of $\mathfrak{g}_2$ was found by locating a generic linear combination of the fourteen basis derivations whose centralizer within $\mathfrak{g}_2$ — computed via the kernel of its adjoint action, using the trace form Tr(XY) as the correct ad-invariant inner product throughout, not a naive Euclidean coordinate inner product — is exactly two-dimensional, confirming $\mathfrak{g}_2$ has rank 2 and yielding an explicit Cartan subalgebra $\{H_1,H_2\}$. The adjoint action of H_1 on the full 14-dimensional space was diagonalized: two eigenvalues are exactly zero (H_1, H_2 themselves), and the remaining twelve form six conjugate pairs $\pm i\omega$, extracted as explicit root vectors $(\omega_1(H),\omega_2(H)) \in \mathbb{R}^2$.

Result. The twelve root vectors found this way split into exactly two length classes in ratio exactly $\sqrt{3}$ (1.7320516…), and every pairwise angle among the six positive roots takes one of exactly five values: 30°, 60°, 90°, 120°, 150° — the complete G_2 signature, matching Section 14.1's geometric root system exactly.

What this is, and what it is not. This is a code-correctness check, not a mathematical discovery and not a prediction: the outcome was already guaranteed by classification theory before any computation ran, so a mismatch could only have indicated a bug in this document's own scripts, not a fact about the world genuinely at risk of turning out otherwise. Its value is accordingly narrow but real — an independent numerical cross-check, using a third method sharing no code with either original construction, that both of this document's own implementations of "G_2" are computing the object they claim to — and it should not be read as more than that.

19.8. Gauge Links: The Kinetic Trace Form Is τ-Invariant; a Raw Wilson Loop Is Not

With τ established as a genuine Lie algebra automorphism (Section 19.7), the natural next question is what this implies for an actual D_4_B lattice gauge theory. A first SO(8) Wilson plaquette was constructed on the D_4 lattice (four independent link variables $U_i = \exp(A_i)$, $A_i \in so(8)$ built from the L, R, $\mathfrak{g}_2$ generators, at two representative root directions), and confirmed to satisfy the standard local gauge invariance expected of any lattice gauge theory: the plaquette trace was unchanged to machine precision (deviation $\approx 6\times10^{-14}$) under independent random gauge transformations at all four corner sites — a basic sanity check on the construction, not yet a test of triality.

Two distinct invariants, tested separately. Replacing every generator A_i in this configuration by its τ-image and rebuilding the plaquette gives a materially different trace (1.914 before, −1.101 after — a genuine, not small, difference). This is the expected outcome, not a failure: the raw Wilson loop trace is a highly nonlinear function of the gauge potentials (a product of matrix exponentials), and bracket-preservation of τ at the Lie algebra level supplies no a priori reason for such a nonlinear quantity, evaluated on an arbitrary field configuration, to be preserved.

The correct invariant to test is the bilinear trace form Tr(XY) that underlies any standard Yang–Mills kinetic term (Tr(F_μν F^μν)) — precisely because this quantity is built linearly from Lie algebra elements. This was verified directly on the same field configurations:

$\mathrm{Tr}(A_1 \cdot A_2) = -1.228197$ and $\mathrm{Tr}(\tau(A_1) \cdot \tau(A_2)) = -1.228197$ (deviation $\approx 4 \times 10^{-16}$); $\mathrm{Tr}(A_1{}^b \cdot A_2{}^b)$ matched to similar precision.

Status. The standard Yang–Mills kinetic term for a D_4_B lattice gauge theory built from this construction is confirmed τ-invariant. A raw Wilson loop trace for an arbitrary field configuration is not, and is not expected to be. What triality implies for matter fields coupled to it in different representations is resolved in Section 19.9. A complete multi-term gauge action, and matter fields in two representations coupled to it, are built and verified consistent in Section 19.10.

19.9. Matter Fields in Different Representations: A Complete, Verified Prescription

Section 19.8 left open what triality implies for matter fields coupled to a D_4_B gauge connection in different representations (8_v, 8_s, 8_c), and left the bracket-preservation check of Section 19.7 as a 10-pair sample rather than exhaustive. Both are resolved directly below.

Exhaustive bracket preservation. The sampled check of Section 19.7 is extended to all $C(28,2) = 378$ pairs of basis generators, not merely ten random ones. The result is unchanged in kind but now complete: exact zero deviation across every single pair, with no exceptions. This has since been repeated in fully exact rational arithmetic, with no floating point or numerical linear algebra anywhere in the chain: the 14-dimensional $\mathfrak{g}_2$ basis was found by solving the defining derivation equations directly over the rationals (confirming a nullspace of exactly dimension 14, not assumed from an SVD cutoff), τ was rebuilt from this exact basis, and all 378 brackets were recomputed and compared symbolically. Every one was found exactly zero, with no exceptions.

The correct prescription, verified exactly. Iterating the automorphism τ itself does not give the correct matter-field transformation, since it answers a different question at the finite group level than the one needed. Instead, for any gauge field value $X \in so(8)$, solving the defining local-triality functional equation $X(x)y + x \cdot Y(y) = Z(xy)$ directly and simultaneously for the pair (Y, Z) — the same linear system used throughout Section 19.5, with a unique solution once antisymmetry is imposed — gives the correct, self-consistent triple (X, X_+, X_-). This triple was verified to satisfy the finite, exponentiated relation

$$\exp(tX)(x) \cdot \exp(tX_+)(y) = \exp(tX_-)(x \cdot y)$$

to machine precision (10^{-13}–10^{-15}) at every tested value of t from 0.001 to 5.0, for three independent test cases — not merely to first order in t, as the Lie-algebra-level checks of Section 19.5 established, but as an exact statement about the finite group elements themselves.

What this establishes. For any D_4_B gauge field configuration, the transformation law for matter in 8_s and 8_c under a given gauge transformation is completely determined by the vector-representation gauge field alone, via a direct linear solve — not by iterating the automorphism τ, which answers the separate question of how the Lie algebra itself relabels under triality (relevant to Section 19.8's gauge-action test). Both quantities are now separately and correctly characterized.

19.10. A Complete Coupled Gauge–Matter Lattice Action, Verified Consistent

Section 19.8 tested τ-invariance of a single bilinear pair; Section 19.9 established how a specific gauge transformation acts on matter in different representations. Both are combined here into a genuine multi-term action with matter, closing the "complete gauge action" item left open at the end of Section 19.8.

A multi-plaquette gauge action. The single-pair trace-form test of Section 19.8 is extended to a genuine sum over six independent plaquettes (distinct pairs among the 24 root directions of D_4), S_gauge = Σ Tr(A_{r_1}·A_{r_2}), each term built from an independently-sampled so(8)-valued potential per link direction. The total action is confirmed τ-invariant to machine precision (deviation ≈ 1.3×10^{-15}) — not merely a single term, as before, but the full sum. This is further extended to a comprehensive set covering every one of the 264 non-parallel pairs among the 24 root directions (44 times the size of the six-plaquette test), confirmed τ-invariant to the same machine precision (deviation ≈ 3×10^{-14}, or 2×10^{-15} relative to the action's own magnitude) — invariance of the gauge sector is not an artifact of a small hand-picked sample.

Matter in two representations, coupled to one gauge field, both gauge-invariant. For a single link with gauge potential A, its local-triality partner A_+ (Section 19.9) was used to build two independent hopping terms — ψ_V(x)^T·exp(A)·ψ_V(x+r) for matter in 8_v, and ψ_S(x)^T·exp(A_+)·ψ_S(x+r) for matter in 8_s — sharing the same underlying gauge field A. Under a local gauge transformation with generator ξ (and its own triality partner ξ_+, obtained the same way), both matter fields and both links were transformed according to the standard non-abelian prescription (ψ → g·ψ, U → g(x)·U·g(x+r)^T), using g = exp(ξ) for the 8_v sector and g = exp(ξ_+) for the 8_s sector. Both hopping terms were confirmed invariant under this transformation to machine precision (deviations ≈ 2×10^{-16} and 9×10^{-16} respectively), with completely independent, generic gauge parameters at the two lattice sites — a genuine test of gauge covariance, not merely a global symmetry check. This has since been extended to all three representations simultaneously (Section 19.12).

Status. This demonstrates a complete, internally consistent lattice construction: a multi-term gauge action, now confirmed across a comprehensive set of plaquette pairs rather than a small sample, and matter fields — extended in Section 19.12 to all three representations — all coupled to a single underlying D_4_B gauge field, with every piece independently verified gauge-invariant. A uniform background gauge field's compatibility with the momentum-space doubler structure is addressed in Section 19.15; a fully dynamical (fluctuating, path-integral) treatment remains beyond this section.

19.11. An Explicit Spatial Lattice Path for the Three-Generation Domain Wall

Section 19.2 established the endpoint values of the mass function separating each pair of triality-related generations (−2 to +2, identical in magnitude across all three walls, forced by T's orthogonality) but left the explicit spatial profile connecting them, and the extension beyond D_4_B, as open. The endpoint construction is completed here.

Construction. Three domain walls are placed at positions s = 0, D, 2D along a single lattice direction (D = 10w, chosen large enough that consecutive walls are well-separated), with

the local mass profile near each wall given by the same tanh form verified in Section 19.1: $M_k(s) = 2\cdot\tanh((s-s_k)/w)$, matching the exact ∓ 2 endpoints established algebraically in Section 19.2.

Verification. The normalization integral for the localized zero mode at each of the three walls was computed independently: all three give exactly 1.333333 (the closed-form value $\int\cosh(s)^{-4}ds = 4/3$ for this mass amplitude), with no discrepancy between them, confirming three genuinely independent, equally-localized, normalizable chiral zero modes — one per generation transition. The mass profile was confirmed to saturate to within 0.01% of its asymptotic ±2 value already at the midpoint between consecutive walls, confirming the three walls do not interfere with one another at the chosen separation.

Status. This gives the domain-wall construction of Sections 19.1–19.2 an explicit, verified spatial realization connecting all three generations, rather than only algebraic endpoint values. The underlying orthogonality proof has since been extended exactly to both order-3 directions of D_4_B's triality structure (Section 19.2); extending it to other root systems, not attempted here, remains the one item from Section 19.2 still open.

19.12. A Genuine Unified Action S = S_gauge + S_Dirac, Verified Gauge-Invariant

Sections 19.8 and 19.10 tested gauge invariance and τ-invariance as separate properties on separate quantities. Here the two are combined into a single object and tested together on an explicit four-site lattice patch.

A first attempt, corrected. Using the open-link bilinear Tr(A_{link1}·A_{link2}) as the gauge term failed local gauge invariance, since that quantity is not a gauge-invariant observable at all; the construction below uses a genuine closed plaquette instead.

The correct construction. Replacing the gauge term with a genuine closed plaquette — the same four-link Wilson loop trace verified gauge-invariant in Section 19.8 — and combining it with four matter hopping terms $\psi(x)^T\cdot U\cdot\psi(x')$ sharing the plaquette's own links gives S = S_gauge + S_matter on an explicit four-site patch (sites x_0, x_1, x_2, x_3 connected along two distinct D_4 root directions). Under independent, generic local gauge transformations at all four sites, S_gauge, S_matter, and their sum S_total were each confirmed invariant to machine precision (deviations 3.6×10^{-15}, 2.5×10^{-14}, and 2.8×10^{-14} respectively).

Extension to all three representations. The construction above covers 8_v and 8_s. It has since been extended to include 8_c as well: for each of the four links, its own local-triality partners (A_+, A_-) were computed once and used to build the corresponding link matrices in all three representations simultaneously, and matter hopping terms were built for 8_v, 8_s, and 8_c together on the same four-site patch, sharing the same underlying gauge field per link and the same underlying gauge-transformation parameter per site (with its own three representation-specific images). Under the same kind of independent, generic local gauge transformation at all four sites, the gauge term, all three matter terms individually, and their combined total were each confirmed invariant to machine precision (deviations ranging from 4.4×10^{-16} to 2.5×10^{-14}). This is the complete picture across all three representations that Sections 19.9–19.10 established piece by piece.

Status and scope. This is the first genuinely unified gauge-plus-matter action in this framework covering all three representations of D_4_B — closer to the S = S_lattice + S_Dirac + S_Yukawa goal named in Section 20 than anything constructed previously. It demonstrates gauge invariance on an explicit finite patch, not the spectral or doubling properties of the full periodic lattice (Section 19.4's momentum-space K(k), W(k) analysis). Matter here is still treated as a pure internal-index field; combining this construction explicitly with the spacetime Clifford (spinor) structure of Section 19.1 is carried out in Section 19.13.

19.13. Combining the Spacetime Spinor Structure with the Internal Gauge Structure

Section 19.12's status paragraph named an explicit gap: matter was treated as a pure internal-index field, without the spacetime Clifford structure of Section 19.1 built in. That gap is closed here.

Explicit Cl(4) gamma matrices. A standard real 4×4 Euclidean representation of the spacetime Clifford algebra was constructed via Pauli-matrix tensor products, and the defining relation $\{\gamma_\mu,\gamma_\nu\} = 2\delta_{\mu\nu}$ was confirmed to hold with exactly zero deviation across all sixteen (μ,ν) pairs — an exact algebraic identity, not a floating-point approximation, since every entry involved is an exact integer or ±i.

A single combined link matrix, not two structures bolted together. The matter field at each site was extended to a genuine tensor-product object, $\psi(x) \in \mathbb{C}^4 \otimes \mathbb{R}^8$ (32 real components: 4 spacetime-spinor components times 8 internal-gauge components), and the hopping link along each D_4 root direction r was built as a single combined 32×32 matrix $\Gamma_r = (r\cdot\gamma) \otimes U_r(x)$, contracting the root's own four spacetime components with the gamma matrices and tensoring with the internal SO(8) gauge link — not two independent pieces multiplied together after the fact.

Gauge invariance of the full combined action. On the same four-site plaquette patch used throughout Sections 19.10–19.12, the combined hopping action $\Sigma\, \psi(x)^\dagger\cdot\Gamma_r\cdot\psi(x')$ was tested under independent local INTERNAL gauge transformations at all four sites (g(x) acting only on the internal factor, $\psi(x) \to \psi(x)\cdot g(x)^T$ in the internal index, leaving the spacetime spinor index untouched — the physically correct transformation law, since internal gauge symmetry does not act on spacetime spinor indices). The gauge term, the full combined matter term, and their sum were each confirmed invariant to machine precision (deviations 3.6×10^{-15}, 9.1×10^{-14}, and 8.7×10^{-14} respectively).

Status and scope. This demonstrates that the spacetime spinor structure of Section 19.1 and the internal gauge structure of Sections 19.9–19.12 combine consistently into a single object with the correct transformation properties, rather than being two structures whose compatibility was merely assumed. Whether Section 19.1's chiral zero mode survives this combination is addressed directly below. What remains beyond this section: extending the construction to all three internal representations simultaneously within this combined 32-component framework, and — as throughout Part VIII — an actual dynamical treatment.

19.14. Does the Chiral Zero Mode Survive Combination with the Internal Gauge Structure?

Section 19.1's domain-wall construction gives a normalizable chiral zero mode from the mass profile M(s) alone, in the absence of any internal gauge structure. Section 19.13 confirmed the combined object transforms correctly under internal gauge transformations, but left open whether the zero mode itself — its normalizability and its chirality — survives once an internal gauge connection is actually turned on, rather than merely checking that the machinery is well-defined.

Setup. A generic internal so(8) generator A was diagonalized; being real and antisymmetric, its eigenvalues are exactly purely imaginary, $\pm i\lambda$ (four conjugate pairs, confirmed directly). For each eigenvalue, the combined zero-mode equation $f'(s) + (M(s) + i\lambda)f(s) = 0$ was solved directly, giving $f(s) = \exp(-\int M(s')ds') \cdot \exp(-i\lambda s)$ — Section 19.1's real exponential envelope, multiplied by a pure phase.

Result. The normalization integral $\int |f(s)|^2 ds$ was computed independently for all eight internal eigenvalues and found identical to Section 19.1's baseline value in every case (2.000000, matching to machine precision, 0 to 4.44×10^{-16}), since $|\exp(-i\lambda s)|^2 = 1$ regardless of λ — an internal gauge charge contributes only a phase, never affecting the real exponential localization. The combined zero-mode equation was verified to hold directly (residual $\sim10^{-10}$, at the limit of numerical differentiation) rather than merely argued analytically. Chirality was checked explicitly on the full 32-component object: $(\Gamma_5 \otimes I_8)\Psi = (+1)\cdot\Psi$ exactly, with zero deviation, since the internal gauge structure acts only on the tensor factor the chirality operator does not touch.

Significance. This is a real question that could have failed: an internal connection with a nonzero REAL part in its spectrum (not antisymmetric, or coupled asymmetrically to the two chiralities) would generically destroy normalizability or split the two chirality components differently. That this document's construction is built entirely from antisymmetric (compact-type) generators throughout is accordingly not incidental to this result — it is the specific structural reason the zero mode survives intact.

19.15. A Uniform Background Gauge Field in Momentum Space: The Doubler Structure Survives Only Within G_2

Section 19.4's doubler-removal result (Section 19.4) was established for the free (ungauged) theory. This section asks a genuinely kinematic question — not a claim of full gauged lattice dynamics — about a fixed, translation-invariant background internal gauge field: one link matrix U_r per root direction, the same at every lattice site, so that ordinary momentum-space (Fourier) treatment still applies exactly, unlike a fully fluctuating dynamical gauge field, which would require a path integral this document does not construct.

First attempt: a generic background, and an honest failure. Twenty-four independent background link matrices, one per root direction, were built from generic random combinations of the full 28-dimensional so(8) generator set, and the extended Wilson mass term $W(k) \to \Sigma_r(I_8 - U_{r,sym}\cdot\cos(k\cdot r))$ was evaluated at k=0. The minimum eigenvalue was 8.20, not zero: the exact physical zero mode of Section 19.4 does not survive a generic

background gauge dressing. A second attempt, using a single shared (rather than independent) background link for all 24 directions, also failed (minimum eigenvalue 0.62): sharing one link across directions is not by itself sufficient.

The precise condition, found and confirmed. Every derivation $X \in \mathfrak{g}_2$ satisfies $X(1) = 0$ exactly — an elementary algebraic fact (from $X(1\cdot1) = X(1)\cdot1 + 1\cdot X(1) = 2X(1)$, forcing $X(1) = 0$) — so any background gauge field confined entirely to $\mathfrak{g}_2$ leaves the octonion unit direction e_0 exactly invariant: $A\cdot e_0 = 0$ was confirmed to hold to machine precision (2×10^{-16}) for an explicit $\mathfrak{g}_2$-valued background, and $\exp(A)\cdot e_0 = e_0$ followed immediately. Repeating the k=0 eigenvalue computation with this $\mathfrak{g}_2$-only background gives a minimum eigenvalue of exactly 0 (deviation -1.8×10^{-15}), while all fifteen doubler points remain strictly positive (minimum eigenvalue 24 or above) — the complete doubler-removal structure of Section 19.4 survives intact.

What this establishes. The momentum-space kinematic structure of Section 19.4 is compatible with a background internal gauge field if and only if that background lies within the triality-fixed subalgebra $\mathfrak{g}_2$, not the full so(8) — a genuinely non-obvious connection between the doubler-removal mechanism of Part VIII and the G_2 substructure studied independently throughout Part VII and Sections 19.5–19.7b, found by direct computation rather than anticipated in advance. A background with any component in L_Im(O) ⊕ R_Im(O) generically gaps the would-be zero mode instead. Whether this is a hint toward, or merely a kinematic curiosity unrelated to, an eventual dynamical mechanism selecting G_2 as the physically realized gauge symmetry (the open question of Section 14.2) is not established here.

20. A Framework for Dynamics

The Chamseddine–Connes spectral action principle derives the full bosonic Standard Model Lagrangian, including gravity, from the spectrum of a Dirac-type operator on a spectral triple. This construction already possesses closely related raw material — the graph Laplacian and covariance matrix of Part III — though building an actual spectral triple from this data remains unperformed. A spacetime Dirac operator (Clifford algebra Cl(4), matching the four-dimensional geometry of Part III) is tensored with, not merged into, the internal Spin(8) representation space of Part V — resolving an apparent dimensional mismatch between the two roles. No dynamical law has yet been derived from any part of this construction: the present work constructs only the kinematical and symmetry framework onto which a genuine dynamics — an action, a Lagrangian, equations of motion — would need to be built. A more tractable starting point than the full spectral action is a deliberately limited action S = S_lattice + S_Dirac + S_Yukawa, built from only the structures already established in Parts III through VII, tested first for the minimal requirements of consistency — whether it preserves the Spin(8) symmetry of Part V, whether it admits the chiral structure sought in Part VIII, and whether it breaks to anything resembling the Standard Model gauge group — before any attempt to reproduce the full Standard Model Lagrangian.

Part IX — The Search for a Non-Circular Selection Principle

21. Criteria Tested and Their Outcomes

A systematic search for a principle selecting D_4 over E_6 or E_8 without presupposing the answer examined several candidates. Symmetry-efficiency and information-theoretic measures were found underdetermined or undefined. Representation-orbit compression numerically favors D_4 but was found circular, rewarding exactly the property (triality-related multiplicity) already known to distinguish it. Spectral-dimension measures do not discriminate at all, since spectral dimension equals rank for every full-rank lattice examined.

A symmetric-split criterion was tested directly: among the seven ways to split $D_8 = so(16)$ into $D_k \times D_{8-k}$, $k = 4$ is the unique self-dual point where the two factors are isomorphic. Tested against this document's own Maximum Symmetry Principle, this criterion fails outright: direct computation of $\|W(D_k)\| \times \|W(D_{8-k})\|$ shows $k = 4$ gives the smallest product (36,864) among genuine two-factor splits, less than half that of $k = 2$ or $k = 6$ (92,160 each) — the symmetric split is disfavored, not favored, by the framework's own existing symmetry principle.

22. A Clean, Non-Circular Result

A distinct, representation-theoretic question was tested: for which n does $D_n = so(2n)$ admit dimensional coincidence between its vector representation (dimension $2n$) and its half-spin representations (dimension 2^{n-1}) — the precise, necessary condition for any permutation symmetry to exist among representations of different geometric origin. Direct verification confirms $2n = 2^{n-1}$ has exactly one solution in positive integers: $n = 4$ ($2\times4 = 8 = 2^3$); for $n < 4$ the vector representation is larger. That no further coincidence occurs for any $n > 4$ follows by induction: at $n = 5$, $2^4 = 16 > 10 = 2\times5$; assuming $2^{(k-1)} > 2k$ for some $k \geq 5$, it follows that $2^k = 2\cdot2^{(k-1)} > 4k$, and since $k \geq 5$ gives $4k > 2k+2 = 2(k+1)$, we obtain $2^k > 2(k+1)$, extending the inequality to $k+1$ and, by induction, to every $n > 4$.

This gives the clearest non-circular explanation found for why triality can exist only for D_4 among all D_n: triality — an exact permutation among three representations — requires those representations to share a common dimension, and this occurs, among all D_n, uniquely at $n = 4$.

22.1. The Internal D_4 Is Forced, Not Chosen, Once the Spacetime Block Is D_4

A question distinct from the preceding one is worth isolating explicitly: granting that the spacetime block D_{4a} is D_4 (established independently in Parts I–III, by the laminated descent's landmark criteria, with no reference to triality), is the appearance of a second D_4 as the internal block D_4_B — the block actually carrying triality — an additional, independent fact requiring its own justification, or a forced consequence of the first?

It is forced. Borel–de Siebenthal theory, already used in Part IV to establish $D_4 \oplus D_4 \subset E_8$, guarantees that any two maximal-rank subalgebras of E_8 of the same isomorphism type are conjugate under the Weyl group $W(E_8)$. Since Weyl group elements act as orthogonal

transformations preserving the root system and all inner products, this conjugacy extends to orthogonal complements: if D_{4a} is any embedding of D_4 realized as one factor of a $D_4 \oplus D_4$ maximal subalgebra of E_8, its orthogonal complement within $\Phi(E_8)$ is the image, under the same Weyl conjugation carrying D_{4a} to the standard embedding, of the standard construction's second factor — and is therefore of D_4 type as well, not merely in the specific coordinate realization used in Part IV but for any such embedding whatsoever. Direct computation confirms this for a second, independently chosen embedding (using coordinates {0,1,4,5} in place of {0,1,2,3}): the resulting subsystem and its orthogonal complement both reproduce D_4's exact root count and inner-product distribution.

The consequence is a genuine reduction in the number of independent facts this framework must postulate. The appearance of a second D_4 hosting triality is not a further coincidence sitting alongside the spacetime block's being D_4; it is a mathematical consequence of the spacetime block's being D_4 at all, via Borel–de Siebenthal alone, with no reference to triality required to derive it. What remains open is exactly one question, not two: why the laminated descent lands on D_4 in the first place (Part II), not a separate question of why the resulting internal structure happens to support triality as well.

23. Necessary and Sufficient Conditions Within the Classification of Simple Lie Algebras

Sections 21–22 show D_4 possesses properties no competitor shares. A materially stronger claim is available: that the properties themselves, demanded in advance, force D_4 uniquely — not merely among D_n, but among every simple Lie algebra.

Theorem (Necessary and Sufficient Conditions for Triality). A simple Lie algebra $\mathfrak{g}$ admits an outer automorphism of order exactly 3 if and only if $\mathfrak{g}$ is isomorphic to $so(8) = D_4$.

Proof. The outer automorphism group of a simple Lie algebra is isomorphic to the automorphism group of its Dynkin diagram (Humphreys, 1972; Bourbaki, 1968). Checking this across the complete classification: type A_n ($n \geq 2$) has diagram automorphism group of order 2; types B_n, C_n, E_7, E_8, F_4, and G_2 have trivial diagram automorphism group (order 1); type D_n ($n \geq 5$) has order 2; type E_6 has order 2. Every one of these has automorphism group of order at most 2, and a group of order 1 or 2 contains no element of order 3. The single exception is D_4: its diagram — a central node joined to three outer nodes — has automorphism group S_3, of order 6, the unique case in the classification containing elements of order 3. No other simple Lie algebra, of any type or rank, admits an order-3 outer automorphism.

This is the converse direction the earlier sections lacked: triality, stated as a property to be demanded independently of any prior preference, admits no solution anywhere in the classification of simple Lie algebras except D_4. Combined with Section 22's explanation of why this holds within the D_n family specifically (the dimension coincidence $2n = 2^{n-1}$ at $n = 4$), this gives a complete, non-circular passage from a stated requirement to a unique structure.

The limit of this result, stated precisely. This theorem has the logical form {desired property} implies D_4 — the converse of {D_4} implies {desired properties} established elsewhere in this document, exactly the direction of implication a selection principle requires. What it does not supply, and what no mathematical theorem can supply on its own, is a derivation of why the desired property (order-3 triality) should be physically demanded in the first place, rather than some other property entirely. That remains a postulate of this framework, named as such in Part X; no amount of further mathematics converts a postulated physical requirement into a proven one — only mathematics conditional on that requirement, which is what this theorem provides.

23.1. A Synthesized Structural Statement and Its Epistemic Status

The results established across Parts I through IX can be organized into a single conditional statement: if a candidate structure is required to be (i) a rootless, even, unimodular lattice, (ii) reached by the laminated descent of Part II under its stated branch-selection rule, (iii) a point at which non-abelian diagram symmetry first appears within the root-lattice range, and (iv) capable of supporting an order-three outer automorphism (triality), then the unique structure satisfying all four conditions is D_4. This statement is a synthesis of results proved separately above, not a new theorem in its own right, and two of its four conditions — (ii)'s branch-selection rule and (iv)'s triality requirement — remain postulates of this framework rather than derived necessities, exactly as Part X discloses; presenting this synthesis as a single unconditional theorem would overstate what has been shown, and it is not so presented here.

23.2. Future Direction: The Missing Variational Principle

The single most consequential open task identified by this entire investigation is not a new computation but a new kind of result: a variational principle from which the requirement of triality itself — condition (iv) above — would follow, rather than being posited. Sections 21 through 23 show that once triality is demanded, D_4 follows uniquely and non-circularly; they do not show why triality should be demanded rather than any other structural property. Locating such a principle, or establishing that none exists within a stated natural class, is identified here as the clearest remaining path toward converting this framework's central postulate into a derived consequence.

Three candidate physical selection principles were tested against this requirement and found not to yield a non-circular derivation: a spectral/maximal-symmetry-density functional (circular by construction), gauge anomaly cancellation (no discriminating power, by Okubo 1977), and modular invariance of the D_4 level-1 Wess–Zumino–Witten (WZW) model (its only independent physical justification belongs to superstring theory, a framework this document does not adopt). This negative result is recorded for completeness, not concealed.

An interesting structural parallel appears in the work of Gresnigt, Gourlay, and Varma (2023, European Physical Journal C 83:747), where an order-three automorphism acting within a sedenionic (Cayley–Dickson) framework likewise generates a three-generation structure with nontrivial algebraic relations among the resulting sectors — a linear

dependence among the three sets of triality-related operators comparable in character to the partial overlap documented for this document's own three triality sectors in Part VII. Although the underlying algebraic setting (sedenions and the Clifford algebra Cl(6)) is fundamentally different from the D_4-triality construction developed here, the recurrence of an order-three symmetry producing three interrelated, partially dependent sectors suggests that such structures may arise more generally than within a single mathematical framework. Whether this similarity reflects a deeper shared mathematical principle or is simply a structural analogy between unrelated constructions remains an open question, not resolved here. The same work also constructs its generation structure through a direct linear automorphism action on fermionic ladder operators, rather than through a domain-wall or Landau–Ginzburg mechanism; this suggests an alternative direction worthy of investigation for the triality-covariant Yukawa coupling identified as missing in Part VIII, though whether such a construction transfers to the D_4 lattice setting used here has not been established.

23.3. An Independent Corroboration via Global Representation Theory

A separate line of inquiry, starting from Conway's group Co_0 itself rather than from the laminated descent, offers independent support for the strategy this document follows, without altering any of its results. The question addressed is: does Co_0's defining 24-dimensional representation on the Leech lattice, considered globally, already carry a complex or quaternionic structure from which triality-like phenomena might emerge directly — or is passage to a smaller substructure genuinely necessary?

Co_0's 24-dimensional representation is documented as absolutely irreducible (irreducible over the complex numbers, not only over the reals). Since Co_0 is defined as a subgroup of the real orthogonal group $O(24)$ preserving the Leech lattice's positive-definite symmetric inner product, this representation carries an invariant real symmetric bilinear form — the defining signature of real Frobenius–Schur type. Its commutant algebra is therefore $\mathbb{R}$ alone: no complex or quaternionic structure is available at the level of the full Conway group.

This is reported as a consistency check, not a new derivation: it confirms, from an entirely different starting point, that no shortcut exists at the level of Co_0's global representation, and that reaching the complex, triality-supporting structure documented in Parts IV through VII requires descending to a genuine substructure — exactly the route this document takes, through the laminated sequence and the explicit $E_8 \supset D_4 \oplus D_4$ decomposition, rather than through Co_0's representation theory directly. The two approaches are complementary: one confirms a necessary condition (no global shortcut exists) that the other's explicit construction already satisfies by taking the only route available.

A note on Co_0 versus Co_1 is worth adding for precision. That Co_0 itself, rather than the simple group $Co_1 = Co_0/\{\pm 1\}$, is required for this representation reflects a separate fact — the central involution acts as $-I$, not as the identity — and is unrelated to the real-type classification above; the two facts should not be conflated.

23.4. A Complete Conditional Proof: Why D_4, Given the Generation-Counting Hypothesis

Sections 23–23.2 established the mathematics but left one link explicitly unclosed: the theorem shows {order-3 outer automorphism} implies D_4 uniquely, but supplies no reason to demand that property in the first place. This subsection states, as precisely and narrowly as possible, the one condition under which that remaining link closes completely — and is equally explicit about what still does not follow, even then.

The hypothesis, stated as a hypothesis, not derived here. Suppose nature requires a simple Lie algebra admitting a symmetry that partitions its structure into exactly three mutually related sectors, as a mechanism for the three observed fermion generations. This is not a mathematical necessity; it is a physical proposal, with its own independent history (Ramond, 1977, and the subsequent literature reviewed in Section 24) and its own competitors (Sections 21, 23.2 record three alternative selection principles that were tested and failed to close this same gap non-circularly). It is stated here as Hypothesis G (generation-counting): *the fundamental symmetry structure relevant to fermion generations is a simple Lie algebra with an outer automorphism group containing an element of order exactly 3.*

The theorem, already proved. Given Hypothesis G, Section 23's theorem applies directly and without modification: among every simple Lie algebra in the complete Cartan classification — A_n, B_n, C_n, D_n, E_6, E_7, E_8, F_4, G_2, at every rank — exactly one, D_4, admits an outer automorphism of order 3. The proof is the diagram-automorphism-group computation already given: every type except D_4 has diagram automorphism group of order at most 2, which contains no element of order 3; D_4's diagram automorphism group is S_3, order 6, uniquely containing one.

The conclusion, and its exact logical status. Modus ponens applies without gap: Hypothesis G, together with the theorem, yields D_4 as the unique admissible structure — not one candidate among several, not the best of a short list, but the only simple Lie algebra satisfying the stated requirement, full stop. Conditional on Hypothesis G, *why D_4* is answered completely: no further mathematical step, and no appeal to any property beyond the single stated requirement, is needed to exclude every alternative. This is the same logical form as any conditional theorem in mathematics or physics — Noether's theorem does not derive why a system possesses a continuous symmetry, only what follows given that it does — and it is presented with the same status here: complete as a conditional argument, silent on the antecedent.

Why the parallel argument for *why 24* has a different character. Section 3.2 gives the closest available analog for the starting dimension: a hypothesis (there, Hypothesis L — minimality combined with well-definedness) under which dimension 24 is shown, completely and non-circularly, to be the unique admissible choice among every classified dimension. The two conditional proofs are not of identical evidential weight, and this is stated there directly rather than left implicit: Hypothesis G rests on a single hard experimental count (three generations), while Hypothesis L rests on

a methodological principle common in physics but not anchored to one specific independent measurement in the same direct way. Both close their respective selection problems completely on their own stated terms; neither hypothesis is itself derived by anything in this document.

What remains open regardless. This subsection closes the D_4-selection gap completely on one explicit condition; it does not, and does not claim to, establish that condition itself. Hypothesis G is exactly the postulate Section 23.2 already named as this framework's central open task — restated here with maximal precision so that what a reader must grant, to accept D_4's selection as fully forced, is stated as a single explicit sentence rather than left implicit across several sections.

24. Relation to Existing Approaches

This pathway is positioned here relative to other research programs, not to argue priority or superiority over them. Superstring theory, the dominant framework for unification, works in ten dimensions and has no established connection to the Leech lattice or D_4; this document diverges from it, as stated at the outset. Loop quantum gravity and causal set theory build spacetime from combinatorial or relational discrete structures without reference to exceptional lattices, sharing this document's general discrete-to-continuum ambition but not its specific starting object. E_8 Theory (Lisi, 2007) and the broader triality-based research program are the closest relatives to Parts IV through VII of this document, and are cited there directly as the source of the speculative bridge to particle physics, not as independent confirmation of it.

A closer and more precise relative to Part V specifically is the trinification family-symmetry program in supersymmetric grand unification, tracing to Glashow's original $[SU(3)]^3 \rtimes \mathbb{Z}_3 \subset E_6$ proposal (1984) and developed in detail by Camargo-Molina, Morais, Ordell, Pasechnik, and Wessén (2018), who construct a supersymmetric model from the identical classical branching $E_8 \supset E_6 \times SU(3)$ used in Part V of this document, explicitly naming the resulting SU(3) a family symmetry SU(3)_F intended to organize three generations. The mechanism is nonetheless distinct from the one developed here: their SU(3)_F is a continuous (global or gauged) symmetry acting on a family index carried by a single, fixed copy of E_6's 27-plet, whereas the triality mechanism of Part VII is discrete ($\mathbb{Z}_3$) and produces three distinct E_6 root systems related cyclically by the verified triality generator, with the spacetime D_4 block held fixed throughout. Both approaches extract a threefold structure from the same $E_8 \supset E_6 \times SU(3)$ branching; whether they are related beyond this shared starting point, or represent genuinely independent mechanisms for the number three, is not established here and is recorded as a further open question this document does not resolve.

The specific idea that triality underlies three fermion generations — the destination Part VII arrives at — is considerably older and more populated than the comparison above might suggest, and this should be stated plainly rather than left for a reader to discover. It was first proposed by Ramond (1977), and has since been developed by a substantial, active research program: Furey's octonionic construction (2014) and its extension with Hughes to all four division algebras (2025), Gresnigt and collaborators' Cl(8) and sedenion

constructions (cited above and elsewhere in this document), and, in the specific setting of the Leech lattice, work explicitly connecting octonionic Leech-lattice constructions to triality and three generations. What this document arrives at in Part VII — triality organizing three generations — is accordingly not a new destination. What appears to differ is the route taken to it: the dominant approach in this literature constructs the relevant division algebra or Clifford algebra directly and derives triality, and often Leech-lattice-adjacent structure, from its automorphisms; this document instead starts from the classical laminated lattice sequence of Conway and Sloane — a purely geometric descent from the Golay code and Leech lattice through intermediate lattices to D_4 — and introduces octonions only afterward, in Part VIII, to make that already-fixed D_4's triality act explicitly on matter representations. Whether this reversal of order (geometric descent first, octonionic algebra second) has appeared elsewhere was checked by web search on two separate occasions, not by a systematic literature review, and no instance was found either time; this is reported as the outcome of that specific, limited search, not as an established claim of priority.

Framework	Starting object	Dimension	Triality	Dynamics
Lisi's E8 Theory	E8	248	Yes	Partial (Lagrangian proposed, chirality disputed)
Superstring theory	10D worldsheet	10	No	Yes (established action)
This framework	Leech lattice	24	Yes	Incomplete (kinematics and symmetry only, Part VIII)

This table is a positioning device, not a ranking; each framework differs in starting assumptions as much as in results, and "incomplete" is stated here about this document's own dynamics with the same directness applied to every other open item in Part X.

24b. Triality Is Inherited from M_{24}, Not Imposed on D_4 Independently

Every criterion tested in Sections 21–23.2 asked, in effect, why physics should *require* order-3 triality as a new, independent principle — and each failed to answer this non-circularly. A different question is asked here: is the triality generator T *already present*, as a residual symmetry inherited from the Mathieu group M_{24} (the automorphism group of the Golay code this entire construction starts from), rather than something imposed on D_4 in isolation? The answer, established below by explicit construction rather than argued for in general terms, is yes: T is not an independent postulate about D_4, but the trace, at D_4, of a symmetry already present in the Golay code this document starts from. The question this reframes — not eliminates — is named at the end of this section.

Construction. The extended binary Golay code was built via the bordered quadratic-residue-on-eleven-points construction of Section 1 and independently reverified (exact weight distribution, exact self-orthogonality). A genuine, nontrivial automorphism of this code was then found — not recalled from memory or assumed from general theory, but located by direct constraint-satisfaction search against the actual verified code data, exploiting the extreme rigidity of the code's S(5,8,24) Steiner system for fast pruning — and verified against all 4096 codewords, not merely a sample. This automorphism has order divisible by 3; the corresponding power gives a genuine order-3 element, independently verified in full. Its action on all 24 Golay coordinates, lifted directly to an isometry of the Leech lattice (Golay code automorphisms act on the Leech lattice by the same coordinate permutation, since the lattice's defining conditions are manifestly permutation-covariant), has the eigenvalue fingerprint +1 (multiplicity 12), and the two complex conjugate cube roots of unity (multiplicity 6 each) — exactly the signature required for a four-dimensional sub-block to carry T's rotation.

Finding a D_4 block and testing it. A "trio" — a partition of all 24 points into three disjoint octads, each independently a genuine Golay codeword — preserved setwise by this automorphism was located by direct search among the code's 759 verified octads. On the resulting 8-point block, the standard E_8 lattice (D_8 plus the half-integer glue class, automatically invariant under any permutation of its 8 coordinates) was built, and the automorphism's own natural invariant 4-point sub-blocks — each consisting of one complete 3-cycle plus one fixed point, guaranteed setwise invariant by the cycle structure itself, not chosen to fit — were used as candidate D_4 blocks. The induced 4×4 orthogonal matrix on each was computed directly.

Result. Both natural 4-point blocks give a matrix with eigenvalues **exactly** +1, +1, $\cos(120°)+i\cdot\sin(120°)$, $\cos(120°)-i\cdot\sin(120°)$ — T's verified signature to the last decimal, not an approximation. Since an orthogonal matrix in four dimensions is determined up to conjugation entirely by its eigenvalues, this is not merely a resemblance: the rotation induced by this genuine M_{24} element **is** T, up to a change of basis. This was independently reproduced with a second, entirely different order-3 automorphism (found via a separate run of the same search), a different preserved trio, and a different octad block — the same exact eigenvalue match resulted both times, indicating the pattern is not an artifact of one particular element.

Statistical robustness. An order-3 orthogonal 4×4 matrix drawn from an 8-point block has, structurally, only one alternative to the T-matching (1,1,3,3) cycle pattern consistent with a permutation of order 3 on 8 points: (5,1,1,1) (five fixed points, one 3-cycle), which does not naturally split into two T-matching 4-blocks. Three further independent order-3 automorphisms were located (varying the target octad in the search), and every trio each one preserves was found exhaustively, not merely the first: across 18 resulting octad-block instances from two of these three automorphisms (the third preserved no trio among the 759 checked — confirming trio-preservation is a real, nontrivial condition, not automatic), every single instance gave the (1,1,3,3) pattern, with zero exceptions.

Explicit identification, not merely matching eigenvalues. Beyond the eigenvalue argument, an explicit orthogonal change-of-basis matrix P was computed directly — bringing both T

and the M_{24}-derived matrix to the same canonical block form via their own eigenvectors and composing the two canonicalizing transformations — giving $P \cdot T \cdot P^T$ = M_trio verified to machine precision (deviation 4.9×10^{-16}), with P itself confirmed orthogonal to the same precision. This is the same style of explicit identification used in Section 19.6 to connect τ|_Cartan to T, applied here to connect an M_{24} element to T directly.

Closing the coordinate gap. The apparent remaining gap — whether this D_4 block coincides with D_4_B as used in Part IV, rather than merely being Weyl(E_8)-conjugate to it — turns out not to require reconstructing Part II's 24-step laminated descent at all. Part IV's own D_4_B (Sections 8–9) is not built by tracing coordinates through the Leech lattice; it is built directly and abstractly from E_8's own 240 roots, by splitting the eight coordinates into two blocks of four — and since E_8 is the unique even unimodular lattice in eight dimensions, this construction depends only on the algebraic formula, not on which physical coordinates realize it. The identical formula, applied to the trio-derived octad block above, gives the same abstract lattice. An explicit orthogonal matrix Q — a pure coordinate permutation, orthogonal to exactly zero deviation — was constructed sending Part IV's own standard first4/last4 D_{4a}/D_4_B split to the trio-derived split exactly: the transformed root sets are supported precisely on the trio construction's coordinate positions, with the correct root count (24) and norm^2 = 2 confirmed. Composed with the P of the previous paragraph, this gives a single, fully explicit chain: a genuine M_{24} element induces, via Q then P, T exactly on the same D_4_B construction Part IV itself uses — not merely a Weyl(E_8)-conjugate copy of it.

The honest limit of this result. What is established: a genuine, verified element of M_{24} — the symmetry group of the object this entire construction starts from — induces T's exact rotation on Part IV's own D_4_B, via an explicit, fully composed chain of orthogonal matrices, confirmed not merely by matching eigenvalues but by explicit change-of-basis identities throughout, found consistently (zero exceptions) across every independent automorphism and every preserved trio checked. Whether this specific identification has already been noted in the literature on Mathieu and Conway groups was checked only by web search, not a systematic review of that literature, and no instance was found — but this is reported as the outcome of a limited search, not as a claim that the result is new to the field: it follows fairly directly from classical facts (E_8's uniqueness, Borel-de Siebenthal conjugacy) well known to specialists in this area, and may be a standard observation or exercise there even where not published as a standalone claim. What is claimed is only that this document did not previously derive it and now does, with full verification. What remains open regardless is not a coordinate technicality but the deeper question this result relocates rather than resolves: *why* nature should realize M_{24} symmetry — the automorphism group of the Golay code — physically in the first place. This is not answered here; it is the same class of question addressed, and left open, in Sections 21–23.2, now asked one level earlier in the descent.

24c. A Genuine S_3, Not Merely Its Order-3 Subgroup

Out(D_4) is the full symmetric group S_3, order 6 — not merely its order-3 (triality) subgroup, which is all Section 24b established. This raised a direct, stronger question: does

M_{24} contain a genuine order-2 element alongside the order-3 one already found, satisfying S_3's defining relation, and acting compatibly on the same structure?

Finding a genuine S_3. Starting from a freshly-found order-3 automorphism M_3 and its preserved octad block (as in Section 24b), a targeted search located an order-2 automorphism M_2 that swaps the block's two natural 4-point sub-blocks — the same sub-blocks (one 3-cycle plus one fixed point each) used to construct T — rather than preserving each separately. Testing all 24 possible internal correspondences between the two sub-blocks located a correspondence satisfying the full S_3 relation $M_2M_3M_2^{-1} = M_3^{-1}$.

The exact count is three, forming one algebraic orbit. Testing all 24 possible internal correspondences located one satisfying the S_3 relation $M_2M_3M_2^{-1} = M_3^{-1}$; an exhaustive, deterministic re-check (removing an initial per-attempt search time limit, which had silently produced two false negatives) found the true count is three. This is not an unexplained empirical count: since M_3^k commutes with M_3, any solution M_2 implies $M_2M_3^k$ is also a solution for every k, so solutions form orbits of size dividing $|\langle M_3\rangle| = 3$ under right-multiplication — and direct verification confirms the three solutions found are exactly one such orbit, $\{M_2, M_2M_3, M_2M_3^2\}$.

Verified three independent ways. The relation holds at the level of the original 24-point permutations; it holds again at the level of the induced 8×8 matrices on the octad block; and generating the group $\langle M_2,M_3\rangle$ directly as 8×8 matrices gives exactly six distinct elements — a genuine, faithful representation of S_3, not a larger or degenerate group. M_2's eigenvalues (four +1's, four −1's) are exactly the signature of a genuine involution, as required.

This S_3 is not the classical one, and that distinction matters. The stabilizer of a trio in M_{24} has the well-documented classical structure $2^6{:}(L_3(2)\times S_3)$ (Curtis, 1976; Ivanov, 2018), with the quotient S_3 acting by *permuting the three octads of the trio* — a standard fact, not a discovery of this document. The S_3 found here is a genuinely different subgroup: both M_2 and M_3 were checked directly and confirmed to preserve *each* octad of the trio individually, never permuting them. This S_3 therefore lives inside the octad-preserving kernel $2^6{:}L_3(2)$, not in the quotient — a finer, more specific structure than the classical trio-permuting S_3, and a different mathematical object from it, even though both are called "S_3" and both act on the same trio. Whether an S_3 subgroup living inside this kernel, matching D_4's $\mathrm{Out}(D_4)$ via an explicit matrix identity of the kind constructed below, has been documented elsewhere was checked directly by web search — narrower and more targeted than the search reported in Section 24b — and no instance was found. This is reported as the outcome of that specific search, not as an established claim of priority: the octad stabilizer's reduction to A_8 (order 20160) is itself classical, and A_8 contains many S_3 subgroups, so specialists aware of both this fact and D_4's triality independently could plausibly identify the same connection without it having been written up as such.

Closing the gap completely. Relating M_2's action to Part IV's D_{4a}/D_4_B convention required two corrections to an initially naive coordinate map. First, since the S_3 relation is an *inversion* ($M_2M_3M_2^{-1} = M_3^{-1}$, not M_3), M_3 must act as T on one sub-block and T^{-1} — not T again — on the other, not the naively assumed "diagonal (T,T)" picture. Second, the two

sub-blocks' actual relative positions within the octad ([0,1,5,6] and [2,3,4,7] in this instance) are not simply "first four, last four," requiring an explicit coordinate permutation composed with the block-diagonal piece rather than a direct placement.

With both corrections in place, a single explicit orthogonal 8×8 matrix P was constructed with

$P^{\top} M_3 P$ = block_diag(T, T^{-1}) and $P^{\top} M_2 P$ = the pure swap matrix [[0,I_4],[I_4,0]],

both confirmed to hold with deviation at the level of 10^{-16} — machine precision, not approximation. This completes the identification begun in Section 24b: the genuine S_3 subgroup of M_{24} found here — distinct from the classical trio-permuting S_3, as established above — corresponds, via a single explicit and fully verified change of basis, to the standard embedding of Out(D_4)'s complete S_3 into the automorphisms of $D_{4a} \oplus D_4_B$ — a twisted-diagonal action (T on one factor, T^{-1} on the other) for the order-3 part, and a pure factor exchange for the order-2 part — not merely up to abstract isomorphism, but as an explicit, checkable matrix identity, exactly as achieved for the order-3 part alone in Section 24b.

The complete group, closed and verified. S_3 has exactly six elements — three rotations, three reflections — and the orbit of three M_2's established above gives exactly the three reflections $\{M_{20}, M_{20}M_3, M_{20}M_3^2\}$. The full 6×6 multiplication table was built directly from these six genuine M_{24} elements (identity, M_3, M_3^2, and the three M_2's) and checked exhaustively: all 36 products land exactly within the same six elements, with zero closure failures. This is not two generators satisfying one relation, extrapolated to a group by assumption — it is a complete, closed, six-element group, verified product by product, matching S_3's abstract structure exactly.

24d. The Full Trio: Six Copies of T, Not One

M_3's cycle type (Section 24b) has six fixed points and six 3-cycles across all 24 Golay coordinates — but only two of each lie within the single octad block used to construct T. This raised a direct question: where are the rest, and do they reveal a larger pattern?

The preserved block extends to a preserved trio. The remaining four fixed points and four 3-cycles were found to combine into two further octads, disjoint from the first and from each other — a complete trio, all three blocks preserved (not merely permuted among each other, as in Section 24e below) by the same single element M_3.

Six copies of T, not one. Each of the three octads splits, by the same construction used throughout this section, into two natural 4-point sub-blocks (one 3-cycle plus one fixed point). All six sub-blocks were checked directly: every one gives eigenvalues matching T's signature exactly, with no exceptions. The total eigenvalue count across all six — twelve +1's, six ω's, six $\bar{\omega}$'s — matches the full 24-dimensional fingerprint established in Section 24b precisely, confirming this accounts for M_3's entire action.

An honest reading of what this is. This is not an independent new finding so much as a direct consequence of the statistical robustness already established in Section 24b: since every fixed-point-plus-3-cycle pairing found there matched T's signature with zero exceptions, and M_3's own cycle type contains exactly six such pairings, this result was effectively

already implied. The count of six sub-blocks does coincide numerically with $|\mathrm{Out}(D_4)| = 6$, giving the arithmetic identity $24 = 6 \times 4 = |\mathrm{Out}(D_4)| \times \mathrm{rank}(D_4)$ — noted here for its elegance, but explicitly *not* claimed as an explanation for why D_4 (rank 4) or the Golay code (dimension 24) are the specific structures this document is built from: the "six" originates in this particular M_{24} element's own cycle type, a fact about M_{24}'s internal structure, not from an independent argument connecting it to D_4's outer automorphism group. Presenting this as progress toward answering "why 4" or "why 24" would overstate what has actually been shown.

The full S_3, not merely M_3, extends across the trio — through two distinct geometric mechanisms. Section 24c's order-2 partner M_2 was checked directly on the other two octads as well, not merely the one used to construct it. M_2 preserves all three octads, and the abstract relation $M_2 M_3 M_2^{-1} = M_3^{-1}$ holds exactly on each of the resulting three pairs of 4-point sub-blocks — verified as an 8×8 matrix identity on every octad, with zero exceptions. What is genuinely striking is *how* this happens: on the first octad, M_2 realizes the relation exactly as in Section 24c, by exchanging the two sub-blocks entirely (a factor swap). On the other two octads, M_2 instead realizes the identical abstract relation by preserving each sub-block separately, acting within each as a direct involution with eigenvalues exactly $(1,-1,1,1)$ — a genuine order-2 element of $\mathrm{Out}(D_4)$ itself, not a factor exchange. Both mechanisms are confirmed to satisfy the same S_3 relation exactly. This shows the single M_{24} element pair (M_2, M_3) realizes one consistent abstract group relation through two qualitatively different geometric actions depending on which part of the 24-dimensional representation is examined — a richer structure than either mechanism alone would suggest, though, as throughout this section, it neither resolves nor is claimed to bear on why 24 or why 4 are the specific numbers this document's descent is built from.

24e. A Second, Structurally Distinct Order-3 Symmetry in M_{24}

Every order-3 automorphism found above (four independent instances, eighteen preserved-trio-block instances, all in Section 24b) shared the same cycle type on the 24 Golay coordinates: six fixed points and six 3-cycles. M_{24} is known to contain two distinct conjugacy classes of order-3 elements. This raised a direct question not yet examined: does the second class exist in this search, and does it behave differently with respect to the trio/D_4 construction used throughout this section?

Finding the second class. Searching for automorphisms that fix individual octad blocks in place — the approach used throughout Section 24b — consistently returned only the first cycle type. Searching instead for automorphisms that permute a trio's three octad blocks among each other, rather than fixing each in place, located an order-6 element with cycle type (6,6,6,6): four six-cycles, no fixed points. Its square, independently verified against all 4096 codewords, is a genuine order-3 automorphism with cycle type (3,3,3,3,3,3,3,3): eight three-cycles, zero fixed points — the second class.

A clean structural distinction, confirmed exactly, not a coincidence. This second class cannot preserve any single octad block: since none of its cycles has length 1, any block it preserved setwise would need to be built entirely from 3-cycles, and 8 is not divisible by 3. This was confirmed directly, not merely argued: exactly 0 of the 759 octads are preserved setwise by

this element, matching the divisibility argument precisely. Instead, this element was found to permute an entire trio's three blocks among each other — each block mapping to a different one, cycling all three as wholes, rather than acting within any single block, as verified directly on an explicit trio.

What this does and does not establish. The two conjugacy classes act on the Golay/Leech structure in geometrically distinct ways: one rotates within a single generation-like block (the class used throughout Section 24b to construct T), the other cycles three such blocks as wholes — arguably a closer structural match, in form, to an intuitive picture of "cycling three generations" than the first class is, though no claim is made here that this bears on that question. This is reported as a clean structural fact about M_{24}, confirmed by direct computation and explained by an exact divisibility argument, not as progress toward resolving any open question named elsewhere in this document: it does not bear on why M_{24} symmetry should be physically realized, and it does not alter the identification of T established via the first class in the body of Section 24b.

Original Contributions of This Work

The individual mathematical objects employed throughout — the Leech lattice, the Golay code, the E_8 root system, D_4 triality, the Gursey–Ramond–Sikivie $E_6/SO(10)$ embedding, the CHSW breaking mechanism — are all classical or previously established elsewhere, as stated in the Introduction. What this document adds, beyond assembling them into one connected pathway, is grouped below by kind, so a reader can distinguish at a glance what is proved, what is computed, what is conceptual framing, and what bears on physical consistency.

Mathematical (proven results)

- the verification, by four independent geometric tests (norm ratio, short-root hexagon angles, long-short angle set, Cartan matrix entry), that this construction's full triality-fixed subalgebra is the exceptional Lie algebra G_2, not merely the SU(3) subsystem of its long roots — identifying the classical origin of the $1/\sqrt{3}$ constant found elsewhere (Part VII, Section 14.1);
- a genuine, independently-found order-3 automorphism of the Golay code (a real Mathieu group M_{24} element, located by direct constraint-satisfaction search rather than assumed or recalled), shown by a fully explicit, composed chain of orthogonal matrices — not merely matching eigenvalues — to induce the triality generator T exactly on Part IV's own D_4_B itself, not merely a conjugate copy of it, reproduced with zero exceptions across four independent automorphisms and eighteen preserved-trio block instances: triality on D_4 is the trace of the Golay code's own M_{24} symmetry, not an independent postulate, with the deeper question this relocates — why nature should realize M_{24} symmetry physically — stated explicitly (Part IX, Section 24b). Whether this specific identification is already known to specialists in Mathieu and Conway group theory — plausible, since it follows fairly directly from classical facts (E_8's uniqueness, Borel-de Siebenthal conjugacy) — was not established by a systematic literature search; "original" here means new to this document, not a verified claim of priority in the field;
- the identification of a conditional structural consequence — six additional color-triplet/antitriplet vector-boson states — following from G_2's standard branching under SU(3), explicitly noted as not practically falsifiable absent a stated symmetry-breaking scale; of its two original assumptions, the first (that SU(3)_color is G_2's maximal SU(3) subgroup) is shown to be a labeling freedom rather than an extra assumption, via a verified Weyl(D_4)-conjugacy check across all 96 candidates, leaving only one genuinely open physical assumption (Part VII, Sections 14.2–14.3);
- the mathematical proof that a cyclic group of order 3 admits no fourth orbit, and the general closure-identity proof that no single fixed mass-coupling direction can support a sign change across all three cyclic triality transitions (Part VII–VIII);
- a Necessary-and-Sufficient-Conditions theorem establishing D_4 as the unique simple Lie algebra admitting an order-three outer automorphism, together with the representation-dimension identity $2n = 2^{n-1}$ explaining this uniqueness within the D_n family (Part IX);

- the Half-E_8 Partition theorem: the union of the three triality generations is exactly half of E_8's roots (Part VII, Section 17.2);
- the identification, from published lattice-classification data, that dimension 24 is not merely the first dimension admitting a rootless even unimodular lattice but the last dimension (among those classified, up to 48) where such a lattice is also unique, with the number of rootless competitors at dimension 32 exceeding ten million, formalized into a complete conditional proof (given an explicit minimality-and-uniqueness hypothesis, stated as such) that 24 is the unique admissible starting dimension among every one classified (Part I, Sections 3.1–3.2);
- the proof, via Weyl-conjugacy under Borel–de Siebenthal, that the internal D_4 hosting triality is a forced consequence of the spacetime block being D_4, not an independent coincidence (Part IX, Section 22.1);
- a complete proof, by exhaustive case analysis over all 18 non-invariant roots of D_4_B, that consecutive triality-related generators are mutually orthogonal ($\langle\alpha,T\alpha\rangle=0$) (Part VIII, Section 19.2);
- the identification that the tetrahedral angles of Section 17.3 and the ⅓:⅔ splitting ratio underlying the orthogonality proof of Section 19.2 are a single geometric fact, not two independent results (Part VIII, Section 19.2).

Computational (machine-verified constructions)

- an explicit, machine-verified computation of the $E_8 \supset D_4 \oplus D_4$ root-system decomposition and its associated (8,8)-bimodule structure (Part IV);
- the direct construction and verification of a triality-invariant SU(3) subalgebra, unique among the 96 candidates tested, cross-checked by two independent computational routes (Part VII, Section 14);
- the explicit construction of three cyclically related E_6×SU(3) breakings, and the complete resolution of their overlap structure into a symmetric 24+16+16+16 decomposition, together with a representation-theoretic derivation of the value 16 from SU(3) branching (Part VII, Sections 16–17);
- the identification and computational refutation of the symmetric-split criterion (Part IX, Section 21) and of an external claim that the three-generation structure follows automatically from $\mathbb{Z}_3$ symmetry alone (Part VIII, Section 19.1);
- the discovery that the three distinct types of 16-root sector within the triality partition share an identical internal inner-product signature, matching the structural pattern of a minuscule spinor representation (Part VII, Section 17.1);
- the discovery that the triality-invariant SU(3) plane relates to all 96 candidate SU(3) planes through exactly two nonzero principal angles, both exact tetrahedral constants, with the full basis-change matrix verified orthogonal with exact algebraic entries (Part VII, Section 17.3);
- an explicit, triality-covariant Wilson term on D_4's 24-root system, verified invariant under T to machine precision, that removes all fifteen lattice doublers (identified via the correct dual lattice D_4^*, not the naive $\mathbb{Z}^4$ coordinate lattice) while preserving a

single exact chiral zero mode at the physical point, closing the tree-level part of Section 19.3's open discretization problem (Part VIII, Section 19.4);

- an explicit octonionic construction of the triality automorphism on so(8), covering all three components (L_Im(O), $\mathfrak{g}_2$, R_Im(O)) accounting for all 28 dimensions, including a from-scratch verification of the Clifford algebra Cl(7) and the relevant Moufang identity underlying it (Part VIII, Section 19.5);
- the discovery that τ restricted to the Cartan subalgebra has an eigenvalue structure (+1, +1, and a genuine 120° complex-conjugate pair) matching the root-space generator T exactly, with $\tau^3 = +I$ confirmed directly, completed by an explicit verified orthogonal matrix P with $P \cdot T \cdot P^T = \tau|_Cartan$ to machine precision (Part VIII, Section 19.6);
- a complete verification that τ is a genuine Lie algebra automorphism of so(8): an exact orthogonal decomposition so(8) = L_Im(O) ⊕ R_Im(O) ⊕ $\mathfrak{g}_2$, an explicit 28×28 matrix representation of τ, and exact bracket preservation confirmed across ten random pairs spanning all three components (Part VIII, Section 19.7);
- an independent numerical cross-check, sharing no code with either original construction, confirming that this document's own implementations of Section 19.5's octonion-derivation $\mathfrak{g}_2$ and Section 14.1's geometric G_2 root system compute the same object — a code-correctness check rather than a new mathematical result, since the identification is already guaranteed by the classification of simple Lie algebras of dimension 14 (Part VIII, Section 19.7b).

Physical (consistency and construction)

- an exact, closed-form continuum solution of the domain-wall chiral zero mode, exponentially localized and of exact algebraic chirality, realized directly with this document's own Cl(4) Dirac matrices (Part VIII, Section 19.1);
- the verification that D_4/Spin(8)'s three fundamental representations are of real Frobenius–Schur type — the same feature underlying the Distler–Garibaldi obstruction against E_8 — together with the demonstration that this construction's specific use of triality avoids that obstruction (Part V, Section 11.1);
- the self-referential wall construction achieving an exact, equal-magnitude sign change across all three triality transitions using no coupling beyond the construction's own generators (Part VIII, Section 19.2);
- a triality-covariant lattice realization of the domain-wall chiral zero mode at tree level, extending the continuum result of Section 19.1 to a discretization compatible with the Nielsen–Ninomiya theorem (Part VIII, Section 19.4);
- confirmation that the standard Yang-Mills kinetic trace form for a first concrete SO(8) lattice gauge configuration on the D_4 lattice is τ-invariant, together with the precise, honest identification of what is not expected to be invariant (a raw Wilson loop trace) and why (Part VIII, Section 19.8);
- the complete, verified prescription for how matter fields in 8_v, 8_s, and 8_c transform consistently under the same gauge transformation, including an initially incorrect approach that was tested, shown to fail at the finite group level, and replaced with a

correct method verified exactly (not merely to leading order) across multiple test cases (Part VIII, Section 19.9);

- a complete multi-plaquette gauge action verified τ-invariant, together with matter fields in two representations coupled to a shared gauge field and independently verified locally gauge-invariant under fully independent gauge parameters (Part VIII, Section 19.10);
- an explicit spatial construction of all three domain walls connecting the triality-related generations of Part VII, with independently verified, exactly matching normalization at each wall (Part VIII, Section 19.11);
- the first genuinely unified gauge-plus-matter action in this framework, combining a closed-plaquette gauge term with matter hopping in all three representations (8_v, 8_s, 8_c) on an explicit lattice patch, with an initial non-gauge-invariant construction tested, found to fail, and corrected, and the resulting action verified fully gauge-invariant for every piece and their combined total (Part VIII, Section 19.12);
- the explicit combination of the spacetime Clifford (spinor) structure of Section 19.1 with the internal gauge structure into a single 32-component matter field and a single combined link matrix, rather than two independently assumed-compatible structures, with the combined action verified gauge-invariant, and Section 19.1's chiral zero mode independently confirmed to survive the combination intact for every internal gauge charge tested (Part VIII, Sections 19.13–19.14).

Conceptual (framework and methodology)

- the Maximum Symmetry Selection Hypothesis, named and explicitly distinguished from a proof, governing branch selection in the laminated descent (Part II);
- the negative result that three independently motivated physical selection principles (a spectral/entropy functional, gauge anomaly cancellation, and modular invariance of the D_4 level-1 WZW model) were tested against the missing variational principle of Section 23.2 and each failed to yield a non-circular derivation, recorded rather than concealed (Part IX, Section 23.2).

This grouping is provided to let a reader distinguish, at a glance, this document's own contribution from the established mathematics it builds upon.

Part X — Status

25. What Is Established, What Is Postulated, What Remains Open

Established:

- The Golay code and Leech lattice's defining properties; Leech's uniqueness under the stated criteria.
- The physical reading of dimension 24 via modular forms and the Schellekens vertex-operator-algebra classification (residual FLM conjecture named).
- Loss of lattice uniqueness beyond dimension 24 (King), confirming 24 as the last dimension where rootlessness selects a single lattice.
- The laminated sequence and its landmark identifications.
- Pointed Gromov–Hausdorff convergence of D_4 to Euclidean $\mathbb{R}^4$, via the diffusive/covariance metric, not graph distance.
- $E_8 \supset D_4 \oplus D_4$ and the $E_6 \times SU(3)$ breaking (CHSW, 1985), both confirmed computationally.
- The full explicit chain to one Standard Model generation's exact quantum numbers via $E_7 \to E_6 \to 27 \to SO(10)$, independently reproduced on a second octad.
- Independent consistency with Harvey–Moore's heterotic construction, sharpened to an explicit verified Weyl(E_8) isometry.
- The triality-invariant SU(3) subalgebra of D_4_B (unique among 96 candidates), satisfying the Nielsen–Ninomiya and Distler–Garibaldi obstructions.
- The three-fold cyclic triality-sector structure of Part VII (Section 16–17): overlap pattern fully resolved (24+16+16+16); all three 16-root sectors share one internal inner-product signature matching a minuscule representation; an exact isomorphism to Part V's SO(10) spinor tested and rejected (signatures differ); the triality-invariant SU(3) plane relates to all 96 candidates via two exact tetrahedral principal angles, with a verified exact basis-change matrix.
- The full triality-fixed subalgebra is G_2, not SU(3) alone (four independent geometric tests) — and identifying SU(3)_color with G_2's maximal SU(3) subgroup is a verified labeling freedom, not an extra assumption, since all 96 candidates are mutually Weyl(D_4)-conjugate (Section 14.3).
- The three generations' union covers exactly half of E_8's roots, with D_4_B entirely excluded; D_4/Spin(8)'s shared real-type chirality obstruction is avoided by this construction's specific use of triality; the internal D_4 hosting triality is forced by Weyl-conjugacy once the spacetime block is D_4, reducing two apparently independent facts to one.
- The dimension-matching identity $2n = 2^{n-1}$ and the Necessary-and-Sufficient-Conditions theorem: D_4 is the unique simple Lie algebra admitting an order-3 outer automorphism $\Leftrightarrow$ n = 4.
- **A genuine, independently-verified order-3 automorphism of the Golay code (a real element of M_{24}, found by direct constraint-satisfaction search, not recalled or assumed) induces T exactly — via a fully explicit, composed chain**

of orthogonal matrices, not merely matching eigenvalues — on Part IV's own D_4_B, not merely a Weyl(E_8)-conjugate copy of it, since Part IV's D_4_B depends only on E_8's defining algebraic formula, not on any specific coordinate route through the laminated descent. Reproduced with zero exceptions across every independent automorphism (four total) and every preserved trio (eighteen block instances) checked. Triality on D_4 is accordingly not an independent postulate but the trace, at D_4, of the Golay code's own M_{24} symmetry; why nature should realize M_{24} symmetry physically remains open (Section 24b).

- **M_{24} contains a genuine S_3 (order 6), not merely the order-3 subgroup used above: an order-2 element satisfying S_3's full defining relation with the order-3 element was found and verified three independent ways (permutation level, 8×8 matrix level, and direct generation of a 6-element group). This S_3 is distinct from the classical trio-permuting S_3 documented in M_{24}'s well-known trio-stabilizer structure 2^6:($L_3(2)$×S_3) (Curtis, 1976; Ivanov, 2018): both elements found here preserve every octad of the trio individually, confirmed directly, placing this S_3 inside the octad-preserving kernel instead. Among the 24 possible internal correspondences, exactly three satisfy the S_3 relation — proved, not merely observed, to form one complete algebraic orbit {M_2, M_2M_3, $M_2M_3^2$} under right-multiplication by ⟨M_3⟩, correcting an earlier time-limited search that had reported only one. A single explicit orthogonal matrix was constructed relating both elements simultaneously to their standard D_{4a}/D_4_B counterparts exactly (deviation ~10^{-16}): a twisted-diagonal (T, T^{-1}) action for the order-3 part and a pure factor exchange for the order-2 part — completing, not merely extending, the identification of Section 24b (Section 24c).**
- **The preserved octad block of Section 24b extends to a complete preserved trio: all six natural 4-point sub-blocks across the three octads match T's exact signature, with zero exceptions, accounting for M_{24}'s entire eigenvalue fingerprint (twelve +1's, six ω's, six $\bar{\omega}$'s). The full S_3 of Section 24c, not merely M_3 alone, was verified to act coherently across all three octads — through two distinct geometric mechanisms confirmed exactly: a factor exchange on one octad, and a direct order-2 involution within each sub-block (eigenvalues exactly 1,−1,1,1) on the other two, both satisfying the identical S_3 relation. Explicitly noted as a consequence of Section 24b's established robustness rather than an independent finding, and explicitly not claimed to bear on why 24 or why 4 are the specific numbers this document is built from (Section 24d).**
- **M_{24}'s second conjugacy class of order-3 elements, distinct from the one used above, is shown to act geometrically differently: it cannot preserve any single octad block (confirmed exactly, 0 of 759, and explained by an exact divisibility argument — 8 is not divisible by 3), and instead cycles an entire trio's three blocks among each other (Section 24e).**
- An exact, closed-form continuum chiral zero mode via real Cl(4) matrices — a genuine kink of a $\mathbb{Z}_3$-symmetric Landau–Ginzburg potential; exhaustive proof of consecutive-generator orthogonality for all 18 non-invariant D_4_B roots, giving a

self-referential three-wall construction with exact, equal sign changes and no external coupling (a single fixed mass-coupling direction was separately shown to support a sign change across only two of the three transitions, refuting the strongest automatic-three-generations claim).

- **A triality-covariant Wilson term on D_4's 24 roots (Section 19.4), removing all fifteen lattice doublers via the correct dual lattice D_4^* while preserving one exact chiral zero mode; an explicit octonionic construction of τ on so(8) across all 28 dimensions, matching T's eigenvalue structure exactly via an explicit change-of-basis matrix, verified a genuine Lie algebra automorphism through exhaustive (378-pair) bracket preservation, since repeated in fully exact rational arithmetic from a $\mathfrak{g}_2$ basis found by direct exact solution rather than numerical SVD, with all 378 brackets exactly zero (Sections 19.5–19.7, 19.9); and an independent numerical cross-check, sharing no code with either construction, that this document's own implementations of Section 19.5's octonion-derivation $\mathfrak{g}_2$ and Section 14.1's geometric G_2 root system are computing the object each claims to — a code-correctness check, not a new mathematical fact, since the identification itself is already guaranteed by classification theory (Section 19.7b).**
- **A background internal gauge field in momentum space (Section 19.15): a generic background destroys Section 19.4's exact zero mode, but a background confined to $\mathfrak{g}_2$ preserves it exactly (deviation ~10^{-15}) while all fifteen doublers remain strictly positive — because every $\mathfrak{g}_2$ derivation annihilates the octonion unit exactly, a fact confirmed directly, not merely cited. This is a genuinely non-obvious connection between Part VIII's doubler-removal mechanism and the G_2 substructure of Part VII, found by computation rather than anticipated in advance.**
- **A first SO(8) Wilson plaquette confirmed gauge-invariant, with the Yang-Mills kinetic trace form (unlike the raw Wilson loop) confirmed τ-invariant; the correct prescription for how matter fields in 8_v, 8_s, and 8_c transform under a gauge transformation, replacing an initially incorrect iterated-τ approach; and a gauge action confirmed τ-invariant across a comprehensive 264-pair set, with matter in two representations confirmed gauge-invariant (Sections 19.8–19.10).**
- **An explicit spatial construction of all three domain walls of Section 19.2, with independently verified, exactly equal normalization at each, and the underlying orthogonality proof extended exactly to both order-3 triality directions (Section 19.11).**
- **The first genuinely unified action S = S_gauge + S_Dirac in this framework, covering all three representations and, separately, combined explicitly with an exact real Cl(4) spacetime spinor structure into a single 32-component matter field — every piece independently confirmed gauge-invariant in each case (Sections 19.12–19.13); and confirmation that Section 19.1's chiral zero mode survives this combination intact — identical normalization for every internal gauge charge and exact chirality preservation — because the internal gauge charge, being purely imaginary (antisymmetric generators throughout),**

contributes only a phase, never affecting the real exponential localization (Section 19.14).

Postulated: Two of these are now closed conditionally rather than left as bare assumptions — Section 3.2 shows dimension 24 and rootlessness follow completely given Hypothesis L (minimality-and-uniqueness), and Section 23.4 shows D_4 follows completely given Hypothesis G (generation-counting) — but the hypotheses themselves remain postulates, un-derived by anything in this document: that Hypothesis L's minimality-and-uniqueness principle, specifically, is the criterion nature applies; that the largest-automorphism-group branch is selected where the laminated construction is non-unique; that the physically relevant symmetry structure is the continuous group Spin(8) rather than the finite lattice automorphism group alone; that Hypothesis G's generation-counting principle, specifically, is the criterion nature applies to select D_4.

Open:

- Why the laminated descent lands on D_4 specifically, rather than E_6 or E_8. This is now closed completely, but only conditionally: Section 23.4 shows that a single explicit hypothesis (a simple Lie algebra with an order-3 outer automorphism, motivated by three fermion generations) forces D_4 uniquely among every simple Lie algebra in the classification, via modus ponens on an already-proved theorem — no further mathematics is needed once that one hypothesis is granted. What remains open is the hypothesis itself, not the mathematics conditional on it; Sections 21 and 23.2 record several independent attempts to derive it non-circularly (four criteria tested in Section 21 against the broader D_4-selection question, three further criteria tested in Section 23.2 against the narrower triality-requirement question specifically), all unsuccessful. The parallel question for dimension 24 is closed the same way (Section 3.2): given a stated minimality-and-uniqueness hypothesis, 24 is the unique admissible starting dimension among every one classified — though that hypothesis is of a methodological rather than directly experimental character, a difference stated explicitly there. A different, related question — whether triality is instead inherited from the Mathieu group M_{24}'s symmetry of the Golay code this construction starts from, rather than imposed on D_4 independently — was answered directly: a genuine M_{24} element induces T exactly on Part IV's own D_4_B, via a fully explicit chain of orthogonal matrices, confirmed with zero exceptions across four independent automorphisms. What remains open there is the question this relocates rather than removes: why nature should realize M_{24} symmetry — the Golay code's own automorphism group — physically in the first place (Section 24b).
- Whether the three cyclic triality sectors of Part VII correspond to physical fermion generations, and what role, if any, their pairwise-shared sectors play.
- The gauge-plus-matter construction of Sections 19.4–19.15 now includes: chirality/zero-mode survival under internal gauge dressing (Section 19.14), confirmed unconditionally; and momentum-space compatibility with a uniform background gauge field (Section 19.15), confirmed only when that background is confined to $\mathfrak{g}_2$, not so(8) generally. Still open: a genuinely fluctuating (dynamical, path-integral) treatment, and whether the $\mathfrak{g}_2$ restriction found in Section 19.15 has

any connection to the separate open question of whether G_2 is physically realized as a gauge symmetry (Section 14.2).

- Extending the orthogonality proof underlying Section 19.11's three-wall construction (established for both order-3 directions on D_4_B) to other root systems.
- A spectral-triple construction realizing dynamics on this construction's own data; any determination of physical coupling constants or particle masses.
- Whether G_2 itself, not only its SU(3) subalgebra, is physically realized as a gauge symmetry — the one assumption the conditional six-state prediction of Section 14.2 still depends on.
- Any prediction distinguishing this construction, observationally, from an alternative. The fourth-generation exclusion implicit in Part VII's mechanism (Section 16) is the clearest candidate, contingent on the open items above — and it should be stated plainly that the absence of a fourth light generation is itself an experimental fact already established since the early 1990s, not a new observation this framework anticipates; consistency with it is the minimum any viable framework must satisfy, not evidence for this one.

This document is offered as a rigorously constructed, honestly bounded mathematical and computational exploration — not as a proof that the observed universe is a projection from a twenty-four-dimensional foundation. That claim, should it ever be established, will require experimental evidence no mathematical construction, however precise, can supply on its own. Its principal contribution is not the discovery of new algebraic objects, but a logically connected, computationally reproducible framework linking well-established mathematical structures into a single speculative pathway from the Leech lattice to four-dimensional triality-based particle symmetry.

Stated plainly, rather than left to be assembled from scattered qualifications: as of this version, the framework offers no experimentally distinguishing prediction. The fourth-generation consequence reproduces an already-established fact; the G_2 color-state consequence is not falsifiable in practice without a stated symmetry-breaking scale; and the tetrahedral-angle geometry is a generic mathematical feature, not a signature specific to this construction. What this document establishes is a verified mathematical and computational consistency chain — not, at this stage, a physical theory with testable content in the ordinary sense. Converting it into one is the task that remains.

Appendix — Computational Reproducibility

Every computational claim in this document (E_8's 240-root system and its $D_4 \oplus D_4$ decomposition, the $E_6 \times SU(3)$ breaking, the order-three triality generator and its action on all 96 candidate SU(3) subalgebras, the three-fold cyclic triality-sector structure and its 24+16+16+16 overlap decomposition, and the Weyl(D_4) symmetry checks of Part III) was produced by direct enumeration in Python using standard numerical linear algebra (NumPy), operating on explicit coordinate representations of E_8 and D_4 root systems. No specialized computer-algebra system or unpublished code library was used. The computations are individually short (root generation by explicit coordinate formulas, inner products, orbit enumeration under a constructed linear map) and are described in the text in sufficient detail — generator matrices, coordinate conventions, and selection rules — for independent reimplementation and verification from those descriptions alone. Throughout this document, the term verified refers to this author's own direct computation, reported with sufficient detail for independent reimplementation, rather than to independent third-party verification, which has not been sought or obtained.

The results of Sections 14.1, 14.3, 17.3, and 19.4 through 19.15 are supported by a complete code appendix of 55 Python/NumPy scripts (51 presented as full source code; 4 explicitly-named failed or superseded attempts presented as a short paragraph each, for readability, with full source retained in the archived repository), provided as a separate companion file (also archived on Zenodo alongside this document, under the same DOI given on the title page) rather than embedded in this document (SciPy is used in thirteen — three for numerical integration via scipy.integrate.quad, eleven for matrix exponentials/logarithms via scipy.linalg.expm/logm, with one script using both; SymPy is used in five, for exact symbolic algebra), organized into seventeen folders matching those sections one-to-one, with a README indexing every script against the specific claim it verifies. One script (Section 14.1) confirms every quantity in the G_2 root-system identification — norms, angles, Cartan integers — exactly and symbolically, not merely to floating-point precision. One script (Section 14.3) confirms all 96 candidate SU(3) subalgebras of D_4_B are mutually Weyl(D_4)-conjugate. One script (Section 17.3) confirms the tetrahedral angle value is a generic consequence of any order-three rotation satisfying the same orthogonality condition, by reproducing it on a randomly generated rotation unrelated to D_4. Six scripts (Section 19.4) construct T on D_4's root system, classify its orbits on the correct D_4* dual-lattice doubler points, construct and verify the triality-covariant Wilson mass term, and check normalizability of the resulting domain-wall zero mode. Seven further scripts (Section 19.5) verify the octonion multiplication table, the Cl(7) Clifford relations, the Moufang identity and resulting local triality triple, the infinitesimal triality relation together with the independent $\mathfrak{g}_2$ fixed-subalgebra check, the R_Im(O) triality partners (including, for transparency, the one analogy-based guess that was tested and found false before the correct linear-solve approach was used), and an exact symbolic (not merely numerical) proof of the core linearized identity. Four scripts (Section 19.6) establish the Cartan-subalgebra membership of L_{e1} and R_{e1}, the corrected 4×4 matrix of τ restricted to the Cartan subalgebra, the explicit change-of-basis matrix identifying it with T, and an exact symbolic (sympy) re-derivation of the same 4×4 matrix confirming every entry is exactly $\pm 1/2$ and τ^3 is exactly the identity, with the underlying

linear system's rank verified exactly rather than assumed from floating-point convergence. One script (Section 19.7) constructs the full 28-dimensional matrix representation of τ and tests exact bracket preservation. One further script (Section 19.7b) is a code-correctness cross-check, not a mathematical result: it confirms this document's own separate implementations of Section 14.1's geometric G_2 root system and Section 19.5's octonion-derivation $\mathfrak{g}_2$ compute the same object, an identification already guaranteed by the classification of dimension-14 simple Lie algebras. Two scripts (Section 19.8) build a first concrete SO(8) gauge configuration on the D_4 lattice and test which invariances do and do not hold under τ, with explicit basis-rank verification. Five scripts (Section 19.9) extend the bracket test to all 378 pairs, reproduce the initially incorrect iterated-τ approach to matter-field transformation (kept for transparency, since the document names this specific failed attempt), isolate its exact cause, verify the corrected finite-group prescription to machine precision, and recompute the exhaustive 378-pair test independently as part of this section's own verification, with explicit rank verification throughout. Three scripts (Section 19.10) build a six-plaquette gauge action, extend it to a comprehensive 264-pair set covering every non-parallel pair of the 24 root directions, and couple matter in two representations to a shared gauge field, all verified consistent and rank-checked. Two scripts (Section 19.11) construct and verify the explicit three-wall spatial path and extend the orthogonality proof of Section 19.2 exactly to T^2 (the other order-3 triality direction). Three scripts (Section 19.12) build the first genuinely unified gauge-plus-matter action — an initial non-gauge-invariant attempt, kept for transparency, the corrected version using a genuine closed plaquette, and its extension to all three representations (8_v, 8_s, 8_c) simultaneously. Two scripts (Section 19.13) build an exact real Cl(4) representation of the spacetime Clifford algebra and combine it explicitly with the internal gauge structure into a single 32-component matter field, verified gauge-invariant. One script (Section 19.14) confirms Section 19.1's chiral zero mode survives this combination intact for every internal gauge charge tested. One script (Section 19.15) tests a uniform background gauge field in momentum space, finding the doubler-removal structure survives only when the background is confined to $\mathfrak{g}_2$. Six scripts (Section 24b) build and verify the exact Golay code, find a genuine order-3 automorphism (a real M_{24} element) via constraint-satisfaction search, test whether it induces T's exact rotation on a natural D_4 block reached via a preserved trio, check this statistically across multiple independent automorphisms, compute an explicit orthogonal change of basis confirming the identification exactly rather than merely by matching eigenvalues, and close the remaining coordinate-identification gap by an explicit permutation matrix connecting the trio-derived D_4 block to Part IV's own D_4_B construction directly. Five further scripts (Section 24c) locate a genuine order-2 M_{24} element completing a full S_3 with the order-3 element, verified three independent ways; construct a single explicit orthogonal matrix relating both simultaneously to their standard D_{4a}/D_4_B counterparts exactly (deviation $\sim 10^{-16}$); confirm this S_3 is distinct from the classical trio-permuting S_3 documented in the literature; prove deterministically that exactly three of the 24 internal correspondences satisfy the S_3 relation — one complete algebraic orbit under right-multiplication by $\langle M_3 \rangle$, correcting an earlier time-limited-search artifact; and build the complete 6-element S_3 multiplication table from six genuine M_{24} elements, confirming all 36 products close correctly. Two further scripts (Section 24d) confirm the preserved octad block extends to a complete preserved trio with all six natural sub-blocks matching T's signature exactly, and that the full S_3 (not merely M_3) acts

coherently across all three octads through two distinct geometric mechanisms. One script (Section 24e) locates M_{24}'s second conjugacy class of order-3 elements and confirms, exactly, that it cannot preserve any single octad block (0 of 759, matching an exact divisibility argument), cycling an entire trio's three blocks among each other instead. Every script was re-run while preparing this appendix and its output checked digit-for-digit against the numbers quoted in the text.

Most scripts in this appendix rely on floating-point linear algebra (NumPy) rather than exact symbolic computation, since the majority of the underlying quantities — while often algebraically simple — involve linear systems (bracket preservation across a 28-dimensional basis, Weyl-conjugacy classification, gauge-field configurations) for which a full symbolic treatment would require substantially more computational machinery for proportionally limited additional certainty beyond the already-tight numerical tolerances (typically 10^{-13} to 10^{-15}, i.e., at the limit of IEEE double-precision representation, not merely "small"). The two scripts most central to the document's claims about exact rational relationships — the core linearized triality identity of Section 19.5 and the Cartan-restricted matrix of Section 19.6 — have been verified exactly and symbolically as representative confirmations that the floating-point results throughout reflect genuine algebraic identities rather than numerical coincidences.

References

Ahn, Y. H. Leptons and quarks from a discrete flavor symmetry. *Physical Review D* 91, 056005 (2015).

Baez, J. C. The Octonions. *Bulletin of the American Mathematical Society* 39, 145–205 (2002).

Borcherds, R. E. The Leech lattice and other lattices. PhD thesis, University of Cambridge (1984); see also *Journal of Algebra* 111, 133 (1987).

Bourbaki, N. *Groupes et algèbres de Lie*, Chapitres 4–6. Hermann (1968).

Camargo-Molina, J. E., Morais, A. P., Ordell, A., Pasechnik, R., Wessén, J. Scale hierarchies, symmetry breaking and particle spectra in SU(3)-family extended SUSY trinification. *Physical Review D* 101, 075031 (2020).

Candelas, P., Horowitz, G., Strominger, A., Witten, E. Vacuum configurations for superstrings. *Nuclear Physics B* 258, 46 (1985).

Chaudhuri, S., Hockney, G., Lykken, J. Maximally supersymmetric string theories in D < 10. *Physical Review Letters* 75, 2264 (1995).

Conway, J. H., Sloane, N. J. A. *Sphere Packings, Lattices and Groups*, 3rd ed. Springer (1999).

Curtis, R. T. A new combinatorial approach to M_{24}. *Mathematical Proceedings of the Cambridge Philosophical Society* 79, 25 (1976).

D'Elia, M., Delfino, F., Nada, A., Vairo, A. Exceptional thermodynamics: the equation of state of G_2 gauge theory. *Journal of High Energy Physics* 2015, 057 (2015).

Distler, J., Garibaldi, S. There is no 'Theory of Everything' inside E_8. *Communications in Mathematical Physics* 298, 419 (2010).

Frenkel, I., Lepowsky, J., Meurman, A. *Vertex Operator Algebras and the Monster*. Academic Press (1988).

Fukushima, M., Oshima, Y., Takeda, M. *Dirichlet Forms and Symmetric Markov Processes*. De Gruyter (1994).

Furey, C. Generations: three prints, in colour. *Journal of High Energy Physics* 2014, 046 (2014).

Furey, C., Hughes, M. J. Three generations and a trio of trialities. *Physics Letters B* 865, 139473 (2025).

Glashow, S. L. Trinification of all elementary particle forces. In *Fifth Workshop on Grand Unification*, K. Kang, H. Fried, P. Frampton (eds.), World Scientific, 88 (1984).

Gresnigt, N., Gourlay, L., Varma, A. Three generations of colored fermions with S_3 family symmetry from Cayley–Dickson sedenions. *European Physical Journal C* 83, 747 (2023).

Gursey, F., Ramond, P., Sikivie, P. A universal gauge theory model based on E_6. *Physics Letters B* 60, 177 (1976).

Harvey, J. A., Moore, G. W. Conway subgroup symmetric compactifications of heterotic string. *Journal of Physics A* 51, 354001 (2018).

Humphreys, J. E. *Introduction to Lie Algebras and Representation Theory*. Springer (1972).

Ivanov, A. A. *The Mathieu Groups*. Cambridge University Press (2018).

Kac, V. G. *Infinite Dimensional Lie Algebras*, 3rd ed. Cambridge University Press (1990).

Kaplan, D. B. A method for simulating chiral fermions on the lattice. *Physics Letters B* 288, 342 (1992).

King, O. D. A mass formula for unimodular lattices with no roots. *Mathematics of Computation* 72, 839 (2003).

Lisi, A. G. An Exceptionally Simple Theory of Everything. arXiv:0711.0770 (2007).

Mikosz, M. Triality, characteristic classes, D_4 and G_2 singularities. *Journal of Homotopy and Related Structures* 10, 323 (2015).

Nielsen, H. B., Ninomiya, M. Absence of neutrinos on a lattice. *Nuclear Physics B* 185, 20 (1981).

Okubo, S. Gauge groups without triangular anomaly. *Physical Review D* 16, 3528 (1977).

Ramond, P. Introduction to Exceptional Lie Groups and Algebras. Technical report, California Institute of Technology (1977).

Schellekens, A. N. Meromorphic c = 24 conformal field theories. *Communications in Mathematical Physics* 153, 159 (1993).

Slansky, R. Group theory for unified model building. *Physics Reports* 79, 1 (1981).